No Map Is Best
A Synthesis of the Cartographic Distortion Programme

Consolidating Three Volumes on Empirical Regularities, Provable Limits, and Statistical Geometry of Map Projections

The closing volume of a four-part cycle, following Original Mathematical Perspectives on Cartography, Empirical Informational Distortion of Map Projections, *and* Metric Embeddings, Impossibility and Information Geometry in Cartography

Silvano Salvador

Contents

Abstract viii

I Synthesis: Coherence Across Three Volumes 1

1 Introduction 2

1.1 Why a fourth volume . . . 2

1.2 What this volume is, and is not . . . 2

1.3 An operational plan for the remaining chapters . . . 3

1.3.1 Chapter 3 (planned): Regime-specific dominance is not a contradiction of impossibility . . . 3

1.3.2 Chapter 4 (planned): Fisher–Rao clustering as an independent cross-validation of the regional/compromise split . . . 3

1.3.3 Chapter 5 (planned, requires new computation): Do the real IDI numbers respect the theoretical distortion bounds? . . . 4

1.4 A stopping rule, stated before the volume is finished rather than after . . . 4

2 A Reader's Map: Three Volumes, One Question 6

2.1 Volume I — *Original Mathematical Perspectives on Cartography* . . . 6

2.2 Volume II — *Empirical Informational Distortion of Map Projections* . . . 7

2.3 Volume III — *Metric Embeddings, Impossibility and Information Geometry in Cartography* . . . 7

2.4 What this map does not yet do . . . 8

3 Regime-Specific Dominance Is Not a Contradiction of Impossibility 9

3.1 The apparent tension, stated precisely . . . 9

3.2 The two volumes quantify over different things . . . 9

3.3 The pool identity, checked explicitly . . . 10

3.4 Why this is exactly what the impossibility theorem predicts, not merely permits 10

3.5 What this chapter does not establish . . . 11

4 Fisher–Rao Clustering as an Independent Cross-Validation 12

4.1 The question this chapter asks . . . 12

4.2 Method: cutting the existing dendrograms at three clusters . . . 12

4.3 Result: a clean match at every regional regime, and an instructive failure at Global 13
4.4 What this chapter does not establish 14

5 Do the Real IDI Numbers Respect the Theoretical Distortion Bounds? 15
5.1 What can be proved: a CDF sandwich 15
5.2 What cannot be proved: a bound on the JSD itself 16
5.2.1 A direct test: holding $D_\rho(x)$ fixed, varying only the sampling 16
5.3 What this leaves for a future investigation 17

6 A Sampling-Concentration Functional for the JSD Bound 18
6.1 Constructing a concentration functional from local, rather than worst-case, distortion 18
6.2 Testing the two-argument bound against a continuous family of concentrations 19
6.3 The moderate range: a genuine, if empirical, two-argument fit 19
6.4 Where it breaks: extreme centre-concentration, and why 20
6.5 What this means for the question Section 5.3 actually asked 21

7 Tail-Sensitive Functionals at Extreme Concentration 22
7.1 Two candidate tail-sensitive functionals 22
7.2 The first candidate fails before it can be tested: near-collinearity with D 22
7.3 The second candidate is well-conditioned but does not move the extreme case . 24
7.4 What would actually be needed, and why this chapter does not build it 25

8 Closing the Cycle 26
8.1 What the five synthesis chapters actually found 26
8.2 Scope boundary 27
8.3 What "coherence" means here, and what it deliberately does not 27
8.4 A final accounting 28

II A Tail-Sensitive Rescue Statistic 29

9 Scope of Part Two 30
9.1 What Section 7.4 actually asked for 30
9.2 The plan for this Part 30
9.3 Why a boundary-rate statistic avoids both earlier failure modes 31
9.4 The sampling model, restated for this Part 32

10 The Exact Distribution of Local Anisotropy, and a General Boundary Theorem 33

10.1 The exact CDF for LAEA . . . 33
10.2 A boundary theorem for any smooth, strictly increasing profile . . . 34
10.3 Direct verification of the boundary rate . . . 35
10.3.1 The naive direct estimator is biased at accessible sample sizes . . . 35
10.3.2 Convergence as $\varepsilon \to 0$: the asymptotic regime is real, just narrower than the first test's grid . . . 35

11 Extreme-Value Convergence and the κ_3 Statistic **37**
11.1 From the boundary rate to an exponential limit law . . . 37
11.2 Simulation check . . . 37
11.3 Defining κ_3 . . . 38

12 Does κ_3 Rescue the Bound Where $\kappa_2^{\max}$ and $\kappa_2^{(2)}$ Did Not? **40**
12.1 The three-argument fit, moderate range . . . 40
12.2 The out-of-range test at $p = 5$ . . . 41
12.3 Reading the fit coefficients . . . 41

13 Pushed Further: Extreme Concentration and Sample-Size Scaling **43**
13.1 Stress test at $p = 20$ . . . 43
13.2 The improvement factor is remarkably stable across x . . . 43
13.3 Is the residual gap a finite-N artefact? . . . 44
13.4 Reading the overlapping standard errors precisely . . . 44

14 A Second Profile: Azimuthal Equidistant **46**
14.1 Setting up a genuinely different closed form . . . 46
14.2 The two- and three-argument fits, AEQD . . . 46
14.3 The out-of-range and stress tests . . . 47
14.4 A modest gap in D produces a large gap in JSD . . . 48

15 Closing This Part, and the Cycle a Fourth Time **49**
15.1 What this Part actually found . . . 49
15.2 Why the theorem succeeding and the bound failing are not in tension . . . 49
15.3 Scope boundary . . . 50
15.4 Why a three-part answer is the right shape for this question . . . 50
15.5 Final accounting: both parts together . . . 51

16 A General Second-Order Correction to the EVT Mean **52**
16.1 What this chapter takes up . . . 52

16.2 A boundary theorem at second order . . . 52
16.3 Specialization to LAEA: an exact closed form . . . 53
16.4 Specialization to AEQD: the same closed form, to a quantified residual . . . 53
16.5 Convergence with n . . . 54

17 Exact Convergence of κ_3's Mean at Extreme Concentration **56**
17.1 The question as left open . . . 56
17.2 Exact quadrature across the sample-size and concentration grid . . . 56
17.3 What this settles . . . 57
17.4 Exact quadrature versus Monte Carlo: a general lesson for this cycle . . . 57

18 A Third Profile: The Gnomonic Case **59**
18.1 Choice of profile . . . 59
18.2 General boundary theorem and EVT limit . . . 59
18.3 The second-order correction is exact here, not merely approximate . . . 60
18.4 What this chapter does not test . . . 60
18.5 What the two exact cases have in common, stated precisely . . . 60

19 The Gnomonic Case, Completed: κ_3's Practical Value for a Third Profile **62**
19.1 The question left open . . . 62
19.2 In-range fit and the $p = 5$ check . . . 62
19.3 Stress test at $p = 20$ . . . 63
19.4 A pattern worth naming, not a proof . . . 63

20 Is the Gnomonic Residual a Finite-n Artefact Either? **65**
20.1 The question, asked a second time . . . 65
20.2 Exact quadrature across the sample-size and concentration grid . . . 65
20.3 What this settles for the gnomonic case . . . 66
20.4 How the two profiles' stabilization values actually diverge . . . 66

21 Completing the Set: Exact Convergence for AEQD **68**
21.1 The one asymmetry left in this cycle . . . 68
21.2 Exact quadrature across the sample-size and concentration grid . . . 68
21.3 What this settles . . . 69
21.4 AEQD's spread against its two companions, precisely . . . 69

22 The Remaining Classical Azimuthal Profiles **71**
22.1 What is actually left untested . . . 71

22.2 Both remaining profiles are degenerate, for structural reasons 71

III Beyond the Azimuthal Family: A Latitude-Band Sampling Model 74

23 Scope of Part Three 75

23.1 Why the cap model does not transfer 75

23.2 A latitude-band sampling model 75

23.3 The profile to test first 76

23.4 What transfers from the cap model's machinery, and what must be re-derived . 77

24 An Exact CDF and General Boundary Theorem for the Latitude-Band Model 78

24.1 The exact CDF for Equirectangular 78

24.2 A general boundary theorem for the band model 79

24.3 Direct verification, and where the direct estimator is biased 80

24.4 What this opens, and what it does not yet close 81

25 Extreme-Value Convergence and κ_3 Under the Band Model 82

25.1 The EVT bridge transfers without a new proof 82

25.2 Simulation check: Equirectangular 82

25.3 The shortfall depends on p as well as on x 84

26 Does κ_3 Rescue the Bound Under the Band Model? 85

26.1 A JSD pipeline for the band model 85

26.2 The three-argument fit 85

26.3 Out-of-range tests: $p = 5$ and $p = 20$ 86

26.4 What the narrow-range result implies for choosing a rescue statistic in practice 87

27 A Second Cylindrical Profile: Where $h \neq 1$ 88

27.1 Choice of profile, and a substitution disclosed rather than made silently 88

27.2 Exact CDF and general boundary theorem 89

27.3 EVT limit and κ_3 for this profile 90

27.4 Practical κ_3 bound-rescue test 91

28 Part Three: Closing Accounting 92

28.1 The bound-rescue result in full 92

28.2 Why Lambert cylindrical equal-area, and what that choice costs 93

28.3 This Part in relation to Part Two 93

IV The Two-Variable Shear Profile **95**

29 Part Four: The Two-Variable Shear Profile **96**

29.1 The profile: Sinusoidal away from the central meridian 96

29.2 Why a new sampling model is needed, and what it assumes 96

29.3 The reduction that makes this tractable 97

29.4 Why the isolated corner is the structurally new case, not merely a harder instance of the old one 98

30 An Exact CDF for the Two-Variable Shear Profile **99**

30.1 The distribution of a product of two power-law variables 99

30.2 Numerical verification 100

30.3 Reading the two branches together 100

31 A Corner-Type Boundary Theorem and Its EVT Bridge **102**

31.1 Why the boundary rate changes shape here 102

31.2 The corner boundary theorem 102

31.3 Numerical verification: the boundary rate 103

31.4 The EVT bridge: a Weibull, not exponential, limit 104

31.5 Numerical verification: the EVT bridge 104

32 Part Four: Closing Accounting **106**

32.1 Why the corner geometry forces a quadratic rate 106

32.2 This Part's place in the cycle's sequence of sampling models 107

33 Correction: the True Tissot Anisotropy **109**

33.1 Tissot's indicatrix from the map's Jacobian 109

33.2 Sinusoidal's true anisotropy 110

33.3 The corrected exact CDF 111

34 The Corrected Boundary Theorem and EVT Bridge **113**

34.1 The corner rate survives the correction; the constant does not 113

34.2 Numerical verification 113

34.3 The corrected EVT bridge 114

34.4 The convergence rates in full 115

35 Mollweide, and What Independence Costs **116**

35.1 Why Mollweide does not inherit the reduction 116

35.2 The shape of $Q_{-}(\phi)$, and why it is not monotone 117

35.3 What the independence assumption in Definition 29.1 does and does not cost . 117

36 Part Four: Final Closing Accounting **119**

36.1 What is closed 119

36.2 What remains open 119

36.3 How large the correction was, and how tightly it was re-verified 120

37 A Real-Data Check: Georgia and France **122**

37.1 The synthetic corner-shear bound-fit 122

37.2 Two real countries 123

37.3 Why real countries sit in a gentler regime 124

37.4 Extending the grid: does an explicit $\lambda_{\max}$ term fix it? 124

References **127**

Abstract

Three earlier volumes asked, by three different methods, why no single map projection minimizes distortion for every purpose. *Original Mathematical Perspectives on Cartography* established the classical and information-theoretic groundwork, including an Informational Distortion Index (IDI) validated on synthetic grids. *Empirical Informational Distortion of Map Projections* tested that index against real GeoNames settlement data across twenty-one countries and five regions, finding a robust, exception-free empirical regularity: regionally-fitted projections beat whole-Earth compromise projections by two to three orders of magnitude, for a growing pool eventually comprising seventeen projections. *Metric Embeddings, Impossibility and Information Geometry in Cartography* then asked three theoretical questions of that regularity: how much distortion is provably forced on any embedding (Part I); whether the empirical regularity is an accident of the tested pool or a provable impossibility (Part II); and what intrinsic geometry the space of projections carries under the same statistical construction (Part III). This volume asks a fourth, narrower question: whether the three methodologies' answers cohere. Chapter 1 states the scope this volume commits to; Chapter 2 is a purely expository, cross-referenced inventory of what each of the three volumes actually proved, with no new claim of its own. Chapter 3 then shows why Volume II's exception-free regional dominance and Volume III's impossibility theorem are not in tension, and Chapter 4 checks whether an independent, metric-free geometric method recovers the same empirical grouping, finding that it does so cleanly at the regime types Chapter 3's argument identifies as relevant, and only partially or not at all elsewhere — reported exactly as found. Chapter 5 attempts the remaining, hardest synthesis question — whether the theoretical volume's closed-form distortion bounds are quantitatively consistent with the empirical volume's real IDI numbers — and proves a genuine, narrower positive result (a CDF sandwich relating the geodesic- and planar-distance distributions under any bi-Lipschitz distortion bound) before showing, by direct computation, that the stronger bound originally sought does not exist: a fixed distortion bound does not determine the resulting JSD once realistic, non-uniform sampling is allowed. No number is asserted here that has not been reproduced from an executed pipeline, and every open question, including this negative result, is stated as such. Chapter 5's own closing section names one further, specific continuation — a two-argument bound incorporating both $D_\rho(x)$ and a concentration functional κ of the sampling distribution — and Chapter 6 takes it up directly: it constructs κ from the profile's local, rather than worst-case, anisotropy, and finds that the resulting bound holds to within a $3.58\times$ span across a wide, realistic range of sampling concentrations, breaking down by an order of magnitude only at an identified, mechanistically explained extreme (near-total centre-collapse of the sample) outside that range. Chapter 6's own closing section names, without building, two candidate tail-sensitive functionals that might rescue that extreme case, and Chapter 7 builds and tests both directly: one proves numerically unstable, nearly collinear with $D_\rho(x)$ itself and unreliable out of range; the other is well-conditioned but leaves the extreme-case underprediction essentially unchanged, directly confirming rather than merely asserting the mechanism Chapter 6 proposed. Chapter 8 declares the four-volume cartographic distortion cycle, now including this third synthesis question, closed.

Part I

Synthesis: Coherence Across Three Volumes

Chapter 1

Introduction

1.1 Why a fourth volume

Three volumes now exist on the same underlying phenomenon: no single map projection minimizes distortion for every purpose, region, or sampling regime. *Original Mathematical Perspectives on Cartography* built the classical and information-theoretic groundwork and proposed the Informational Distortion Index (IDI), an empirical diagnostic validated on synthetic grids. *Empirical Informational Distortion of Map Projections* took that diagnostic to real GeoNames settlement data across twenty-one countries and five regions, and found the pattern repeatedly, without exception, across a pool that grew from twelve to seventeen projections: projections fitted to a region (Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic, and, in the extended pool, Bonne and Polyconic) beat whole-Earth compromise projections (Robinson, Mollweide, Aitoff, and others) by two to three orders of magnitude. *Metric Embeddings, Impossibility and Information Geometry in Cartography* then asked, of that empirical pattern, three theoretical questions belonging to three different branches of mathematics: how much distortion is provably forced on any embedding of a region of the sphere into the plane (metric embedding theory); whether the empirical no-uniform-winner pattern is a proved impossibility rather than an artefact of the specific projections tested (a no-free-lunch argument); and what intrinsic geometry the space of projections carries once each one is treated as inducing a point on a statistical manifold (information geometry, via the Fisher–Rao metric).

Three books answering three genuinely different questions about one phenomenon leave a fourth question open, which this volume takes up directly: whether the three answers cohere, in the sense of one explaining, constraining, or cross-validating another. The chapters this volume plans (Section 1.3) answer it by checking specific, falsifiable claims of coherence against the actual content of the three volumes.

1.2 What this volume is, and is not

This is not a fourth empirical study extending the sampling regimes, the country pool, or the projection set further; Section 1.4 explains why that particular kind of extension is being declined. It is not a popularization or abridgement of the three existing volumes; a reader who wants their content should read them. It is a synthesis in the narrow, specific sense that follows: a small number of concrete claims of the form *result A from one volume and result B from another are logically or empirically consistent, in this specific way, checkable as follows* — together with, where such a check has not yet been carried out, an explicit statement that it has not, rather than a plausible-sounding paragraph standing in for it.

Chapter 2 is the one chapter of this kind completed so far, and it is deliberately the least ambitious: a factual, cross-referenced inventory of what each of the three volumes actually proved or found, with no synthesis claim of its own, serving as the common reference point every

later chapter (once written) will cite rather than re-summarize. Section 1.3 below lays out, in operational detail, the three substantive synthesis chapters this volume is committed to, what each requires, and which of the three already has everything it needs versus which requires new computation.

1.3 An operational plan for the remaining chapters

Three synthesis chapters are planned. They are listed in increasing order of how much new work each requires, from none to a genuinely new pipeline run. Sections 1.3.1 and 1.3.2 below are now written, as Chapters 3 and 4; each turned out to need a small, disclosed amount of new computation beyond what was anticipated here, noted in place. Section 1.3.3 remains as originally planned, not yet attempted.

1.3.1 Chapter 3 (planned): Regime-specific dominance is not a contradiction of impossibility

Claim to be checked. The empirical volume's headline finding — a fixed small set of regionally-fitted projections dominates a fixed set of compromise projections at every one of twenty-one countries tested — might look, superficially, in tension with the theoretical volume's impossibility theorem, which says no projection is uniformly best. Chapter 2's inventory already makes the resolution visible without further computation: the impossibility theorem's own statement is that no projection dominates *across sampling regimes and region shapes taken adversarially against each other* (its numerical illustration works with four, then five, synthetic regimes — Global, Regional equatorial, Regional mid-latitude, Regional high-latitude, and Arctic — chosen precisely to be mutually adversarial), whereas the empirical volume's twenty-one countries are not adversarial against the regionally-fitted trio (or quintet) in this sense: every one of them is a genuine, bounded region, which is exactly the regime the regionally-fitted projections are constructed for. The empirical finding is a demonstration of the impossibility theorem's own mechanism (different regimes favour different projections) rather than a counterexample to its conclusion (no projection wins in every regime, including whole-Earth ones). **What this chapter needs**: no new computation — only a careful, explicit argument, citing Volume III's impossibility theorem (§ 6–8 there) exact hypotheses and Chapter 2's inventory of the empirical volume's exact regime and country list, laid out side by side. This is the most straightforward of the three planned chapters and could be written next.

1.3.2 Chapter 4 (planned): Fisher–Rao clustering as an independent cross-validation of the regional/compromise split

Claim to be checked. The theoretical volume's Part III found, via unsupervised hierarchical clustering under the Fisher–Rao metric (structurally unrelated to the Jensen–Shannon divergence the IDI itself is built from), a grouping of projections into conic-family and whole-sphere-compromise clusters, recovered independently at all four sampling regimes tested there and stable under resampling. If this grouping lines up with the empirical volume's regionally-fitted-versus-compromise split — Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, and Lambert Conformal Conic landing in one Fisher–Rao cluster, Robinson, Mollweide, Aitoff, and the other compromise projections in another — that is a second, independent confirmation of the same empirical split, obtained by a method that never looks at a single real settlement coordinate. **What this chapter needs**: no new computation either, in the

narrow sense that Volume III's regional-clustering dendrograms (Part III, extended-regimes chapter) already exist and need only be read against the empirical volume's own projection list; it does need a careful check of exactly which projections appear in both the theoretical volume's twelve-to-seventeen-projection pool and the empirical volume's own list, since the two pools were built independently and are not guaranteed to overlap completely — this check has not yet been carried out and is the chapter's first concrete task.

1.3.3 Chapter 5 (planned, requires new computation): Do the real IDI numbers respect the theoretical distortion bounds?

Claim to be checked. The theoretical volume's Part I proves closed-form or tightly bounded global bi-Lipschitz distortion for five classical azimuthal projections as a function of a spherical cap's angular radius $x = r/R$, and a universal lower bound $D_{\min}(x)$ over *all* embeddings. The empirical volume's real, GeoNames-derived IDI values are a different quantity (a Jensen–Shannon divergence between distance distributions, not a bi-Lipschitz distortion ratio), computed on real, non-uniform, irregularly-shaped national point clouds rather than idealized spherical caps. Before any quantitative comparison is meaningful, this chapter would need to establish, explicitly, a bridge between the two quantities — for instance, via Volume III's already-proved local quadratic relation between JSD and Fisher–Rao distance (Part III), combined with a further, not-yet-attempted relation between Fisher–Rao distance and bi-Lipschitz distortion — and only then check whether, say, the Lambert Azimuthal Equal-Area projection's real empirical IDI at a given country's angular size sits consistently relative to $D_{\text{ea}}(x)$ at that country's approximate angular radius. **What this chapter needs**: this is the one planned chapter that requires genuinely new work — both a piece of mathematics (the JSD-to-distortion bridge, not yet proved anywhere in the three existing volumes) and a new computational pipeline (extracting each tested country's approximate angular radius from its real geographic extent, already available in the empirical volume's GeoNames data, and joining it against the theoretical volume's closed-form bounds). No placeholder numbers are given for it here; if undertaken, this chapter would either succeed in establishing a genuine quantitative link or report, exactly as the empirical volume reported its failed scaling-law attempt, that the bridge does not hold cleanly — either outcome would be reported as found, not adjusted to look more conclusive.

1.4 A stopping rule, stated before the volume is finished rather than after

The toponomastic series closed on the explicit ground that a further extension would risk an unbounded sequence of ever-narrower follow-on pools; the theoretical cartography volume's own Part I closed its azimuthal-profile investigation on the explicit ground that its remaining open question (the Global Embedding Optimality Problem, stated formally in that volume's closing Part I chapter) needed a different kind of mathematics than the one carrying the chapters up to that point. This volume commits, in advance rather than in retrospect, to an analogous rule: once Sections 1.3.1 to 1.3.3 are written, the cartographic distortion cycle — four volumes, one empirical index, one impossibility theorem, one differential geometry, and this synthesis — is closed. Genuinely new questions arising from this cycle's own content (for instance, Section 1.3.3's JSD-to-distortion bridge, if it turns out to open a rich vein of its own) are acknowledged as legitimate future work, exactly as the Global Embedding Optimality Problem was in the theoretical volume, but are explicitly not committed to as a fifth cartography

volume by default. Any continuation would need its own, separately stated justification, not the inertia of an already-running series.

Chapter 2

A Reader's Map: Three Volumes, One Question

This chapter makes no synthesis claim. It is a factual, cross-referenced inventory of what each of the three earlier volumes actually established, organized so that Sections 1.3.1 to 1.3.3's planned chapters (Section 1.3) can cite it rather than re-summarize their sources each time. Every claim below is attributed to a specific chapter of a specific volume; nothing here is a new computation or a new argument.

2.1 Volume I — *Original Mathematical Perspectives on Cartography*

This volume's first seven substantive chapters each isolate one classical or less-classical theorem and trace its cartographic consequence, moving from the purely qualitative to the constructive: Gauss's Theorema Egregium, as an absolute obstruction to any locally isometric map (differential geometry); Chebyshev's variational characterization of the optimal conformal projection for a given region (calculus of variations); the Hairy Ball Theorem, as a topological constraint on globally continuous wind-field representations (algebraic topology); angular defect and the discrete Gauss–Bonnet theorem underlying polyhedral projections such as Fuller's Dymaxion map (discrete geometry); the logarithmic-spiral geometry of loxodromes; and the classical characterizations of the gnomonic and stereographic projections as circle-preserving constructions (projective and inversive geometry, respectively). These six chapters are independent of one another and of the volume's later empirical programme; no later volume, including this one, depends on their specific content beyond the general cartographic vocabulary (great-circle distance, projection formulas) they establish.

The volume's final theoretical contribution is the Informational Distortion Index (IDI) itself: a Jensen–Shannon divergence between the distribution of true geodesic distances and the distribution of projected planar distances, computed over a sampled grid, proposed as a distributional complement to pointwise classical distortion measures (Tissot's indicatrix, angular and areal error). This is validated computationally, on synthetic latitude/longitude grids, with the volume's own stated discipline that all such numerical results are the direct, unedited output of an executed pipeline.

A second empirical strand, opened once the IDI existed, tests whether national-scale geodesic-distance-distribution structure correlates with three independent country-level measures: toponymic (place-naming-convention) divergence, two country-shape descriptors (elongation, isoperimetric deficit), and a fourth measure built from real digital-elevation-model terrain data attached to the same twenty-one-country GeoNames records used elsewhere in the cycle. The headline finding of this strand, stated in the volume's own abstract, is a null result honestly reported rather than suppressed: no evidence of a toponymic–geodesic relationship is found, and the volume's closing four-axis correlation chapter finds all four measures (toponymic, two shape descriptors, elevation) to be substantially complementary rather than redundant, together with an explicitly hedged exploratory country typology built from them.

2.2 Volume II — *Empirical Informational Distortion of Map Projections*

This volume takes the IDI, as constructed in Volume I, to real data: populated-place coordinates from the GeoNames gazetteer, for twenty-one countries across five geographically and culturally diverse regions (Mediterranean Europe, the Maghreb, the Caucasus and Anatolia, South Asia, and Atlantic/Western Europe). Twelve projections, already implemented in Volume I, are re-scored against these real, non-uniform, population-concentrated point clouds, region by region.

The volume's central, repeatedly-confirmed empirical regularity is this: at national and regional scale, projections centred and fitted to the region under test — Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, and Lambert Conformal Conic — systematically outperform projections optimized for whole-Earth compromise display — Robinson, Mollweide, Hammer, Winkel Tripel, Van der Grinten, and Aitoff — by two to three orders of magnitude in IDI, with no exception across all twenty-one countries tested. A separate chapter puts this finding through explicit robustness scrutiny (sensitivity to arbitrary pipeline parameters, and a null check against gazetteer size), and a further chapter attempts, but does not achieve, a clean analytic scaling law for how the IDI should behave as a sampled region shrinks to a point: the histogram-based JSD estimator is found to saturate in a combinatorial noise floor before any asymptotic regime is reached, at every one of six angular scales tested (a 32-fold range from 20° down to 0.625°), and this negative result is reported as a failure rather than adjusted or omitted.

A later extension, added to the same volume without revising any of its original twelve-projection numbers, introduces five further projections (Gall–Peters, Eckert IV, Eckert VI, Bonne, Polyconic) and reruns the full twenty-one-country pipeline once, from scratch, on all seventeen projections together. This does not overturn the original finding but sharpens it: the three-member regionally-fitted trio turns out to be the leading edge of a five-member quintet (adding Bonne and Polyconic), with no new exception introduced anywhere the original study did not already report one. A further robustness pass on this extended quintet finds its occupation of the leading ranks exactly as robust as the original trio's, but its internal ordering measurably less so, and reports two specific correlation questions (including one against the Bonne/Polyconic margin, made possible only by the extension, since neither projection existed in the original pool) as explicitly open rather than resolved.

2.3 Volume III — *Metric Embeddings, Impossibility and Information Geometry in Cartography*

This volume is explicitly framed as an independent companion to Volume I, taking the IDI as a point of departure rather than a further object of empirical testing, and pursues three theoretical questions in three independent parts.

Part I asks how much distortion is provably forced on any embedding of a spherical cap into the Euclidean plane. It proves an explicit chordal (orthographic-projection) distortion benchmark, $\sec(r/R)$, first on radial point pairs and later, via a general theorem developed further into the part, over every pair in the cap; a universal lower bound $D_{\min}(r)$ over *all* possible embeddings, via Ptolemy's inequality; an exact closed-form global distortion $D_{\text{eq}}(r) = (r/R)/\sin(r/R)$ for the azimuthal equidistant projection, narrower than the chordal benchmark; closed forms for three further classical azimuthal projections (equal-area and stereographic, sharing the ex-

act value $\sec^2(r/2R)$; gnomonic, $\sec^2(r/R)$, markedly worse than the other four), unified under a single curvature-sign theorem; and a proof that the azimuthal equidistant projection is a strict local minimum of distortion along every admissible smooth perturbation direction. The part closes not by resolving the remaining global-optimality question — whether $D_{\text{eq}}(r)$ equals the true infimum over all embeddings, azimuthal or not — but by formalizing it into three explicitly named, logically independent sub-problems (a symmetrization step with no known rearrangement-inequality analogue for this kind of functional; a global variational step obstructed by the functional's unestablished convexity; an existence-of-minimizer step) and stating plainly that none of the three is attempted.

Part II proves a formal impossibility theorem for the IDI itself: no single projection minimizes it uniformly across a fixed collection of mutually adversarial sampling regimes. The theorem is illustrated numerically first against Volume II's own twelve-projection pool and four synthetic regimes (Global, Regional equatorial, Regional mid-latitude, Regional high-latitude), by direct, executed enumeration, and then re-illustrated against the seventeen-projection extended pool and a fifth (Arctic) regime, where a *third* distinct regime-wise winner emerges, if anything sharpening rather than softening the impossibility. Both illustrations use synthetic sampling regimes constructed for the purpose, not Volume II's real per-country GeoNames data.

Part III develops the Fisher–Rao geometry of the statistical manifold underlying the IDI's construction: it establishes an exact isometry between the Fisher information metric of the relevant categorical model and a round sphere of constant curvature $1/4$, an exact Fisher–Rao distance formula, a proved local quadratic relation between this distance and the Jensen–Shannon divergence, and a demonstration that Part II's regime-wise IDI-minimizing projections remain the regime-wise Fisher–Rao-minimizing projections as well (a metric-independence check). A further chapter computes full pairwise Fisher–Rao distance matrices and hierarchical clustering at all four (not only the original one) synthetic regimes of Part II's own construction, recovering a consistent grouping of projections into conic-family and whole-sphere-compromise clusters, checked for robustness across three linkage methods and found stable (mean Adjusted Rand Index between 0.83 and 0.99) under 80% grid resampling at every regime tested.

2.4 What this map does not yet do

This chapter has not compared any two volumes' findings against each other; it has only laid out, side by side, what each independently contains. In particular, it has not yet checked whether Volume II's twenty-one real countries fall within the regime types Volume III's Part II treats as adversarial (Section 1.3.1); it has not yet checked whether the specific projections appearing in Volume II's empirical trio or quintet are the same projections Volume III's Part III clusters into its conic-family group, nor confirmed the two volumes' projection pools actually overlap in the way that comparison would require (Section 1.3.2); and it has not attempted any quantitative bridge between Volume III's closed-form distortion bounds and Volume II's real, GeoNames-derived IDI numbers (Section 1.3.3). Those three checks are exactly Section 1.3's remaining operational plan, in the order given there.

Chapter 3

Regime-Specific Dominance Is Not a Contradiction of Impossibility

What this chapter does, and does not, require

Section 1.3.1 committed this chapter to no new computation: only an explicit argument, laid out from the two volumes' own stated hypotheses and data, showing that Volume II's regional empirical dominance and Volume III's impossibility theorem are not merely compatible by coincidence but are, properly read, two descriptions of the same underlying fact. That argument is given below in full.

3.1 The apparent tension, stated precisely

Volume II's headline finding, recalled in Section 2.2, is that a fixed trio of regionally-fitted projections — Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic — beats a fixed set of whole-Earth compromise projections at *every one* of twenty-one countries tested, with no exception. Volume III's Part II impossibility theorem, recalled in Section 2.3, proves that no projection in a twelve-projection pool (identical, name for name, to Volume II's own original pool; Section 3.3 below confirms this explicitly) minimizes the IDI uniformly across a four-regime collection.

Read carelessly, these look like they pull in opposite directions: one volume reports a dominance pattern with literally no exception across twenty-one test cases; the other proves, for the same twelve projections, that uniform dominance is impossible. Section 3.2 shows this is not a tension at all, once each volume's precise domain of quantification is made explicit rather than left implicit.

3.2 The two volumes quantify over different things

Volume III's impossibility theorem, in its own exact statement (recalled in Section 2.3), quantifies existentially over a collection of *sampling regimes* $\mathcal{R}$ deliberately constructed to be mutually adversarial: Global (a full-sphere grid spanning the entire range of latitudes and longitudes), Regional equatorial, Regional mid-latitude, Regional high-latitude, and, in the extended illustration, Arctic. These five regimes are synthetic grids built for the specific purpose of stressing every candidate projection against a spread of latitude coverage, from the whole sphere down to narrow bands at four different positions on it. The theorem's content is that this specific adversarial spread is enough to defeat every candidate in the pool: whichever projection wins at Global (a full-sphere regime) is beaten at at least one of the narrower regional regimes, and vice versa.

Volume II's twenty-one countries, by contrast, are not instances of Volume III's adversarial

regime collection at all. Every one of them is what Volume III's own Part II plausibility remark (recalled in Section 2.3) calls a genuine, bounded region: a national territory, always narrower in geographic extent than a whole-sphere grid, always the kind of compact, latitude-and-longitude-bounded sample that Volume III's own regional regimes (not its Global regime) were built to represent. Volume II never tests any of its three regionally-fitted projections, or any compromise projection, against a Global-type regime at all — its pipeline (Section 2.2) evaluates every projection recentred on each country's own sample centroid, precisely the condition under which Volume III's own plausibility account (Section 2.3) predicts a fitted-centre or fitted-parallel construction should do well, and a whole-sphere compromise construction should not.

Key Result

Volume II's twenty-one-country dominance pattern is not a counterexample to Volume III's impossibility theorem because it does not test the theorem's adversarial condition at all: it tests only one side of the adversarial pair (bounded regional grids), never the other (a full-sphere grid), for any of its twenty-one countries. The theorem's claim is that no projection wins on *both* sides of the adversarial spread; Volume II's finding is that a specific trio wins consistently on *one* side of it. These are different claims about different test conditions, and both can be, and are, true simultaneously.

3.3 The pool identity, checked explicitly

This resolution is strengthened, not merely permitted, by the fact — confirmed directly from both volumes' own text, not assumed — that Volume III's twelve-projection impossibility pool (Equirectangular, Mercator, Sinusoidal, Lambert Azimuthal Equal-Area, Lambert Conformal Conic, Robinson, Winkel Tripel, Mollweide, Aitoff, Albers Equal-Area Conic, Hammer, Van der Grinten) is, name for name, identical to Volume II's own original twelve-projection pool, and Volume III's extended seventeen-projection pool (adding Gall–Peters, Eckert IV, Eckert VI, Bonne, Polyconic) is stated in Volume III's own text to reproduce Volume II's five-projection extension *verbatim*, with none of the five formulas altered. This is not a coincidence requiring further explanation: Volume III explicitly reuses Volume II's projection implementations, by its own account (Section 2.3). It does mean the "same twelve projections" quantified over in both volumes' claims really are the same set, so that Section 3.2's resolution is a statement about the same objects under different regime conditions, not a comparison across two only loosely related pools.

3.4 Why this is exactly what the impossibility theorem predicts, not merely permits

Section 3.2's resolution can be stated more strongly than mere non-contradiction. Volume III's own plausibility remark (Section 2.3) already offers a mechanism for *why* the regime-wise winners split the way the impossibility theorem's own regime-by-regime IDI table shows them splitting: a fitted-centre or fitted-parallel construction (Lambert Azimuthal Equal-Area, and by the same logic Albers Equal-Area Conic and Lambert Conformal Conic) is exact, or nearly so, in a neighbourhood of its own fitted centre or parallels, with distortion growing outward; a compact regional grid sits almost entirely within that low-distortion neighbourhood, while a global grid does not. Volume II's twenty-one countries are, by construction, compact regional

grids recentred on each country's own centroid — precisely the condition under which this mechanism predicts the fitted trio should dominate. Volume II's finding is therefore not an independent empirical fact merely *consistent* with Volume III's theoretical account; it is closer to a direct empirical instantiation, at national scale and on real settlement data rather than synthetic grids, of the specific mechanism Volume III's Part II already proposed — informally, as a plausibility remark, not a further theorem — to explain its own impossibility result. Volume II was written, and its data collected, before Volume III's plausibility remark was; the agreement is not built into either volume's design, in the same sense Volume III's own synthesis chapter (recalled in Section 2.3) reports its Aitoff/LAEA cross-validation as unplanned.

3.5 What this chapter does not establish

This chapter has not checked whether Volume II's specific twenty-one countries, if fed through Volume III's own regime-classification pipeline, would in fact register as instances of Volume III's "Regional" regime type rather than some intermediate case — it has argued from the two volumes' stated methodologies and design intentions that they should, not from a direct computational check assigning each of the twenty-one countries a regime label under Volume III's own criteria. It has also not addressed whether a single country large enough to span a genuinely global range of latitudes (none of Volume II's twenty-one do, but none is a priori excluded from existing) would behave as Volume III's Global regime does; this remains an untested boundary case. Neither gap affects Section 3.2's central resolution, which rests on the two volumes' regimes being different in kind (bounded-region versus full-sphere) as stated in their own text, not on a quantitative regime-classification exercise this chapter has not performed.

Chapter 4

Fisher–Rao Clustering as an Independent Cross-Validation

A small amount of new computation, disclosed

Section 1.3.2 anticipated this chapter would need "no new computation," only a reading of Volume III's existing dendrograms against Volume II's projection list. That turned out to be almost, but not quite, true: Volume III's own chapter reports only the *first three merges* of each regime's dendrogram (Section 2.3), which is not the same thing as a full three-cluster partition, and does not by itself say which cluster Robinson or Van der Grinten — both absent from Volume III's own prose, though present in its pool — end up in. Section 4.2 below closes that gap with the smallest computation that could do it: cutting Volume III's own, already-computed, already-verified pairwise Fisher–Rao distance matrices into three clusters directly (`code/cross_volume_cluster_check.py`), rather than reading off merge order. No new distance matrix, grid, or IDI value is computed anywhere in this chapter; only an existing matrix is cut at a different threshold.

4.1 The question this chapter asks

Volume III's Part III clusters the twelve pool projections by pairwise Fisher–Rao distance — a construction that never once looks at a real settlement coordinate, a real country boundary, or the Jensen–Shannon-divergence-based IDI score any specific projection achieves on real data. Volume II's regionally-fitted trio (Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic) was identified by exactly the opposite route: real GeoNames settlement coordinates, real national boundaries, and the IDI itself. If Volume III's geometry-only clustering places these same three projections in one cluster, separate from Volume II's compromise sextet (Robinson, Mollweide, Hammer, Winkel Tripel, Van der Grinten, Aitoff), that is a genuine independent cross-validation: two methods that share no data and no construction arriving at the same grouping.

4.2 Method: cutting the existing dendrograms at three clusters

Section 3.3 already confirmed the two volumes' twelve-projection pools are identical, name for name, so the comparison is between the same twelve objects under both methods. For each of Volume III's four regimes, `code/cross_volume_cluster_check.py` loads the regime's already-computed 12×12 pairwise Fisher–Rao matrix (`data/regional_pairwise_fr_results.json`, produced by Volume III's own pipeline and not recomputed here), re-runs the identical average-linkage clustering Volume III's own chapter used, and cuts the resulting tree at exactly three clusters (`scipy`'s `fcluster` with `maxclust=3`) rather than reporting merge order. This is a different question from the one

Volume III's own prose answers ("which projections merge earliest") and needed a small script of its own; it does not require, and does not perform, any new distance computation.

4.3 Result: a clean match at every regional regime, and an instructive failure at Global

Table 4.1: Three-cluster partition (average linkage, cut at $k = 3$) of the twelve pool projections, by regime. Bold marks Volume II's regionally-fitted trio; italics mark Volume II's compromise sextet. Source: `code/cross_volume_cluster_check.py`, `data/cross_volume_cluster_results.json`, run on Volume III's own already-verified pairwise Fisher–Rao matrices.

Regime	Three-cluster partition
Global	{Equirectangular}; {Mercator, **LCC**}; {Sinusoidal, **LAEA**, **Albers**, *Robinson*, *Winkel Tripel*, *Mollweide*, *Aitoff*, *Hammer*, *Van der Grinten*} — trio *not* unified: LCC separates with Mercator (Section 2.3's exact-identity finding), while LAEA and Albers fall into one large cluster together with the entire compromise sextet.
Regional mid-latitude	{Equirectangular}; {Sinusoidal, **LAEA**, **LCC**, **Albers**}; {Mercator, *Robinson*, *Winkel Tripel*, *Mollweide*, *Aitoff*, *Hammer*, *Van der Grinten*} — trio unified, and *no* compromise-sextet member shares its cluster.
Regional equatorial	{Equirectangular}; {Mercator, Sinusoidal, **LAEA**, **LCC**, **Albers**, *Aitoff*, *Hammer*, *Van der Grinten*}; {*Robinson*, *Winkel Tripel*, *Mollweide*} — trio unified, but three of six sextet members (Aitoff, Hammer, Van der Grinten) share the trio's cluster; only Robinson, Winkel Tripel, and Mollweide separate cleanly.
Regional high-latitude	{Equirectangular, Mercator, *Van der Grinten*}; {Sinusoidal, **LAEA**, **LCC**, **Albers**}; {*Robinson*, *Winkel Tripel*, *Mollweide*, *Aitoff*, *Hammer*} — trio unified, and five of six sextet members separate cleanly; only Van der Grinten defects, joining Equirectangular and Mercator instead.

Table 4.1 shows a graded, not uniform, result, and it is reported exactly as computed rather than smoothed toward the cleanest-sounding regime. At **Regional mid-latitude**, the match is complete: the trio (Lambert Azimuthal Equal-Area, Lambert Conformal Conic, Albers Equal-Area Conic) forms its own cluster, and every one of the six compromise-sextet projections (Robinson, Mollweide, Hammer, Winkel Tripel, Van der Grinten, Aitoff) — Volume III's own chapter prose named only four of these six explicitly (Section 2.3); Robinson and Van der Grinten's membership is this chapter's own finding — falls into the other substantive cluster, with none crossing over. At **Regional high-latitude**, the same clean separation holds for five of the six sextet members; only Van der Grinten defects to a small residual cluster with Equirectangular and Mercator, an honestly-reported exception rather than a rounding error. At **Regional equatorial**, the separation is markedly weaker: only half the sextet (Robinson, Winkel Tripel, Mollweide) clusters apart from the trio, while the other half (Aitoff, Hammer, Van der Grinten) shares the trio's cluster — this is the same regime Volume III's own prose (Section 2.3) already flagged as the one where "this trio does not cluster as tightly," and the three-cluster cut confirms that flag rather than contradicting it. At **Global**, the three-cluster cut does not recover a trio-versus-compromise split at all: Lambert Conformal Conic separates early with Mercator

(an exact identity Volume III proves elsewhere, Section 2.3), while Lambert Azimuthal Equal-Area and Albers Equal-Area Conic are absorbed into one large cluster together with the *entire* compromise sextet.

Key Result

Fisher–Rao clustering recovers Volume II's regionally-fitted trio as its own cluster, cleanly separated from every member of Volume II's compromise sextet, at the Regional mid-latitude regime completely, at Regional high-latitude with one honest exception (Van der Grinten), and at Regional equatorial only partially (half the sextet). At the Global regime, the three-cluster cut does not separate the trio from the compromise sextet at all. This is not an unqualified cross-validation of Volume II's finding: it is a cross-validation that is strong precisely at the regime types Chapter 3 already argued are the relevant analogue of Volume II's twenty-one bounded countries (Regional, not Global), and weak or absent at the one regime type (Global) Chapter 3 argued Volume II never actually tests. The two chapters' findings support each other: an independent geometric method separates the trio from the sextet exactly where, and to the degree that, Chapter 3's argument says it should.

4.4 What this chapter does not establish

This chapter has cut Volume III's existing dendrograms at $k = 3$ clusters, a natural choice matching the trio-versus-sextet-versus-outliers structure being tested for, but has not checked whether some other number of clusters, or a different linkage rule beyond the average linkage Volume III's own Section 2.3 already found agrees exactly with complete linkage at all four regimes, would tell a materially different story. It has also not explained *why* the Regional equatorial and high-latitude regimes specifically are the two where the separation is imperfect, beyond repeating Volume III's own unexplained observation (Section 2.3) that these are the same two regimes where its clustering signal was already found comparatively fragile by two unrelated diagnostics (linkage-rule sensitivity and resampling stability); this chapter's finding is a third, independent symptom of the same unexplained fragility, not a resolution of it. Finally, this chapter has said nothing about *why* Lambert Conformal Conic separates from its own trio at the Global regime specifically (via the exact Mercator identity), a Global-regime-specific mechanism already proved in Volume III but not one this chapter investigates further.

Chapter 5

Do the Real IDI Numbers Respect the Theoretical Distortion Bounds?

What was attempted, stated before the result

Section 1.3.3 left this chapter open because no mathematical bridge between Volume III Part I's bi-Lipschitz distortion $D_\rho(x)$ and Volume I/II's Jensen–Shannon-divergence-based IDI existed in any of the three source volumes. This chapter attempts to build one directly. It succeeds in proving a genuine, if narrower, positive result (Section 5.1) and then uses that same machinery to show, by explicit computation rather than assertion, that the stronger bridge originally hoped for — a bound taking $D_\rho(x)$ alone to a bound on the JSD — does not exist (Section 5.2). Both outcomes are reported as found. Neither was adjusted to look more or less conclusive than the computation actually shows, in keeping with Volume II's own precedent of reporting its failed scaling-law attempt (Section 2.2) rather than omitting it.

5.1 What can be proved: a CDF sandwich

Fix an azimuthal profile ρ and cap radius x, and write, for a pair of points at colatitudes and azimuths within the cap, $\text{ratio}(p,q) := d_{\text{planar}}(p,q)/d_{\text{geo}}(p,q)$. Volume III Part I's own machinery (recalled in Section 2.3) establishes, for every profile it treats, that this ratio lies in a known interval (m, M) with $M/m = D_\rho(x)$ (for the Lambert Azimuthal Equal-Area profile at radius x, $m = \cos(x/2)$, $M = \sec(x/2)$, $M/m = \sec^2(x/2)$).

Proposition 5.1 (CDF sandwich under bounded ratio)**.** *Let μ_{geo} and μ_{planar} be the distributions of $d_{\text{geo}}(p,q)$ and $d_{\text{planar}}(p,q)/m$ (planar distances rescaled by the ratio's own infimum m) over any distribution of point pairs (p,q) drawn from the cap, and let $F_{\text{geo}}, F_{\text{planar}}$ be their respective CDFs. Then, writing $D = M/m$,*

$$F_{\text{planar}}(t) \;\le\; F_{\text{geo}}(t) \;\le\; F_{\text{planar}}(Dt) \qquad \text{for every } t \ge 0. \tag{5.1}$$

Proof. After rescaling, $\text{ratio}'(p,q) := d_{\text{planar}}(p,q)/(m\, d_{\text{geo}}(p,q)) \in [1, D]$ for every pair, so $d_{\text{geo}}(p,q) \le d_{\text{planar}}(p,q)/m \le D\, d_{\text{geo}}(p,q)$ pointwise, for every pair, regardless of how pairs are sampled. If $d_{\text{planar}}(p,q)/m < t$ then $d_{\text{geo}}(p,q) \le d_{\text{planar}}(p,q)/m < t$, giving $\{d_{\text{planar}}/m < t\} \subseteq \{d_{\text{geo}} < t\}$ and hence $F_{\text{planar}}(t) \le F_{\text{geo}}(t)$. If $d_{\text{geo}}(p,q) < t$ then $d_{\text{planar}}(p,q)/m \le D\, d_{\text{geo}}(p,q) < Dt$, giving $\{d_{\text{geo}} < t\} \subseteq \{d_{\text{planar}}/m < Dt\}$ and hence $F_{\text{geo}}(t) \le F_{\text{planar}}(Dt)$. □

Proposition 5.1 is checked numerically in `code/distortion_jsd_bridge_check.py`: for the Lambert Azimuthal Equal-Area profile at $x = 0.3$ ($D = \sec^2(0.15) = 1.022842$), 4000 uniformly sampled cap points, and 50 threshold values t spanning the observed distance range, both inequalities in Eq. (5.1) hold with zero violations once the planar distances are rescaled by the theoretical infimum $m = \cos(x/2)$ — a rescaling this check confirmed was necessary: an

earlier, uncorrected run of the same check against the *raw*, unrescaled planar distances showed 22 of 50 threshold violations of the first inequality, exactly because the raw ratio for this profile ranges over $(\cos(x/2), \sec(x/2)) \not\subseteq [1, D]$ rather than $[1, D]$ directly; the theorem as stated requires the rescaling, and does not hold without it.

Key Result

A bi-Lipschitz distortion bound $D_\rho(x)$ does imply a genuine, provable constraint relating the geodesic- and planar-distance distributions: their CDFs are sandwiched as in Eq. (5.1), a multiplicative analogue of a Lévy–Prokhorov-type distance bound. This is a real, if narrow, positive answer to part of Section 1.3.3's question.

5.2 What cannot be proved: a bound on the JSD itself

Section 1.3.3's original hope was stronger than Proposition 5.1: a bound taking $D_\rho(x)$ to a bound on the Jensen–Shannon divergence the IDI is actually built from, since that is the quantity Volume I/II's real pipeline computes, not a CDF sandwich. This section shows, by direct construction rather than by a failed proof attempt alone, that no such bound can exist using $D_\rho(x)$ by itself.

The reason is structural. Proposition 5.1's sandwich constrains *where mass can be*, multiplicatively, along the distance axis; the JSD, computed (as in Volume I's own methodology, Section 2.1) from a *histogram* of sampled pairwise distances with a fixed bin width, depends additionally on how much probability mass is concentrated near any given distance scale — information the sandwich alone does not supply. Two point-pair distributions can satisfy identical CDF sandwich bounds while having arbitrarily different JSD, if one concentrates mass where the sandwich's multiplicative slack matters more than the other does.

5.2.1 A direct test: holding $D_\rho(x)$ fixed, varying only the sampling

Table 5.1 fixes the Lambert Azimuthal Equal-Area profile and a cap radius x — hence fixes $D_{\text{ea}}(x)$ exactly, the one quantity Section 1.3.3 hoped would predict the resulting JSD — and varies only the spatial distribution of the sampled points within the cap: uniform, concentrated toward the centre (where distortion is smallest), or concentrated toward the boundary (where it is largest). This directly mimics the real difference between an idealized uniform sampling grid and a real country's actual, geographically uneven settlement pattern — precisely the difference between Volume I's synthetic-grid IDI validation and Volume II's real GeoNames-based measurement.

Table 5.1: Mean JSD (bits, 10 replicates of 250 points each, 40-bin histograms) between the geodesic- and LAEA-planar-distance distributions, at three cap radii, for three spatial sampling concentrations. $D_{\text{ea}}(x)$ is identical across all three sampling columns at a given radius, by construction — it depends only on x and the profile, not on how points are sampled within the cap. Source: `code/distortion_jsd_bridge_check.py`, `data/distortion_jsd_bridge_results.json`.

x	$D_{\text{ea}}(x)$	JSD, uniform	JSD, centre-conc.	JSD, boundary-conc.	max/min ratio
0.15	1.005646	5.70×10^{-6}	4.45×10^{-6}	1.32×10^{-5}	2.96×
0.30	1.022842	3.41×10^{-5}	2.20×10^{-5}	1.26×10^{-4}	5.73×
0.50	1.065199	1.77×10^{-4}	1.04×10^{-4}	1.01×10^{-3}	9.68×

At every radius tested, boundary-concentrated sampling produces substantially higher JSD than centre-concentrated sampling — unsurprising given the underlying geometry (Section 2.3: the Lambert Azimuthal Equal-Area profile's ratio departs furthest from 1 near the cap boundary), but not something $D_\rho(x)$ alone, a single worst-case scalar over the whole cap, encodes; $D_\rho(x)$ is completely blind to *where* within the cap that worst case is realized relative to where the sample actually concentrates. The sensitivity is not small and grows with x: a nearly $3\times$ spread at $x = 0.15$ widens to almost $10\times$ at $x = 0.5$, at fixed $D_{\text{ea}}(x)$ throughout.

Key Result

$D_\rho(x)$ does not determine the JSD-based IDI, even for a fixed profile and fixed cap radius: holding both fixed and varying only the spatial concentration of the sampled points changes the resulting JSD by a factor of 3–$10\times$ across the radii tested. No bound of the form "JSD $\leq g(D_\rho(x))$" for any function g can be correct in general, since the same $D_\rho(x)$ value here manifests three different JSD values depending on sampling alone. Section 1.3.3's hoped-for quantitative comparison between Volume III's closed-form distortion bounds and Volume II's real IDI numbers is accordingly not achievable through $D_\rho(x)$ directly: real countries' settlement distributions are exactly the kind of non-uniform sampling this section shows breaks any such bound, and are moreover unique to each country rather than drawn from one of three clean concentration archetypes, compounding the problem rather than resolving it.

5.3 What this leaves for a future investigation

This negative result is specific to bounding JSD by $D_\rho(x)$ *alone*; it does not rule out a bound that additionally incorporates information about the sampling distribution's own concentration — for instance, a bound of the form JSD $\leq g(D_\rho(x), \kappa)$ where κ is some concentration functional of the real point distribution (its own analogue of a density or an entropy relative to the cap's area measure). Constructing such a κ and proving a genuine two-argument bound is not attempted here; Section 5.2.1's experiment used three qualitatively distinct, hand-chosen concentration archetypes precisely because they were enough to falsify the one-argument bound Section 1.3.3 hoped for, not because they exhaust the space of concentration functionals a more refined bound might use. Real countries' settlement distributions, being neither uniform, nor centre-concentrated, nor boundary-concentrated in any clean sense, would in any case require their own concentration functional computed directly from GeoNames data — a further empirical step this chapter has not taken and does not commit this monograph to taking. Section 2.4's original question — whether Volume III's distortion bounds are quantitatively consistent with Volume II's real IDI numbers — is accordingly answered here in the only honest way currently available: not through $D_\rho(x)$ alone, which this chapter has shown cannot work, and not yet through any two-argument refinement, which this chapter has not constructed.

Chapter 6

A Sampling-Concentration Functional for the JSD Bound

Scope

Section 5.3 named a specific continuation: a two-argument bound JSD $\leq g(D_\rho(x), \kappa)$ for some concentration functional κ of the sampling distribution. This chapter constructs a concrete κ (Section 6.1), tests it against Chapter 5's sampling-sensitivity design extended to a continuous family of concentrations (Section 6.2), finds that it substantially — not completely — rescues the bound over a realistic range (Section 6.3), and finds an identifiable failure at an extreme of that range (Section 6.4).

6.1 Constructing a concentration functional from local, rather than worst-case, distortion

Section 5.1's $D_\rho(x)$ is a single worst-case scalar: the ratio M/m between the largest and smallest local stretch the profile achieves anywhere in the cap. It says nothing about *where* in the cap that worst case occurs relative to where the sampled points actually are, which is exactly what Section 5.2 showed matters. The natural fix is a functional of the sampling distribution's own weight on the cap's local distortion, not a second worst-case scalar.

For the Lambert Azimuthal Equal-Area profile, $\rho(\phi) = 2R\sin(\phi/2)$ for colatitude ϕ from the cap centre. Its radial and tangential scale factors are, respectively,

$$h_\phi(\phi) = \frac{1}{R}\frac{d\rho}{d\phi} = \cos(\phi/2), \qquad h_\theta(\phi) = \frac{\rho}{R\sin\phi} = \frac{1}{\cos(\phi/2)}, \tag{6.1}$$

whose product is 1 everywhere, as an equal-area projection requires. Their ratio,

$$\delta(\phi) := \frac{h_\theta(\phi)}{h_\phi(\phi)} = \sec^2(\phi/2), \tag{6.2}$$

is the profile's *local* anisotropy at colatitude ϕ — and $\delta(x) = \sec^2(x/2) = D_{\text{ea}}(x)$ exactly: the cap-wide worst-case distortion used throughout Chapter 5 is simply δ evaluated at the boundary, where it is largest, since δ increases monotonically with ϕ.

Definition 6.1 (Sampling-weighted local distortion)**.** For a sampling distribution of points over the cap with colatitude marginal density $f(\phi)$,

$$\kappa := \mathbb{E}_{\phi\sim f}[\delta(\phi)] = \int_0^x \delta(\phi)\, f(\phi)\, d\phi. \tag{6.3}$$

κ is the *typical* local distortion a sampled pair actually experiences, as opposed to $D_\rho(x)$, the distortion the *worst-placed* pair would experience. By construction $1 \leq \kappa \leq D_\rho(x)$ for any sampling distribution supported on the cap, with $\kappa \to 1$ as sampling concentrates at the centre

and $\kappa \to D_\rho(x)$ as it concentrates at the boundary. Unlike $D_\rho(x)$, κ is computable directly from a sample: $\hat{\kappa} = \frac{1}{n}\sum_{i=1}^{n} \delta(\phi_i)$ over the n sampled points, with no further model of the sampling distribution required. This is the concentration functional Section 5.3 left unconstructed; `code/kappa_bound_check.py` implements Eq. (6.3) exactly as $\hat{\kappa}$.

6.2 Testing the two-argument bound against a continuous family of concentrations

Chapter 5's own sampling-sensitivity experiment (Section 5.2.1) used three hand-chosen archetypes — uniform, centre-concentrated, boundary-concentrated — precisely enough to falsify the one-argument bound, but not enough to fit or falsify a two-argument one. This chapter replaces the three archetypes with a one-parameter family: sample the cap radial coordinate as $u = U^p$ for $U \sim \text{Uniform}(0, 1)$ and exponent $p > 0$, so that $p < 1$ concentrates points toward the boundary, $p = 1$ recovers uniform sampling over the cap's area measure, and $p > 1$ concentrates points toward the centre — the same three regimes Chapter 5 tested, now continuously parametrized rather than sampled at three isolated points.

Table 6.1 reports D, $\hat{\kappa}$, and the mean JSD (over eight replicates of 1500 points, 40-bin histograms) at the same three cap radii Chapter 5 used, for nine values of p spanning strong boundary-concentration ($p = 0.2$) to strong centre-concentration ($p = 5$).

Table 6.1: D, sampling-weighted local distortion $\hat{\kappa}$ (Eq. (6.3)), and mean JSD (bits, 8 replicates of 1500 points, 40-bin histograms) across nine concentration exponents p at three cap radii. Source: `code/kappa_bound_check.py`, `data/kappa_bound_results.json`.

x	p	D	$\hat{\kappa}$	JSD
0.15	0.20	1.005646	1.004692	1.05×10^{-5}
0.15	1.00	1.005646	1.002801	9.93×10^{-7}
0.15	3.00	1.005646	1.001401	4.85×10^{-7}
0.15	5.00	1.005646	1.000937	8.59×10^{-7}
0.30	0.20	1.022842	1.018935	1.70×10^{-4}
0.30	1.00	1.022842	1.011269	1.40×10^{-5}
0.30	3.00	1.022842	1.005628	5.65×10^{-6}
0.30	5.00	1.022842	1.003760	1.39×10^{-5}
0.50	0.20	1.065199	1.053723	1.32×10^{-3}
0.50	1.00	1.065199	1.031728	1.07×10^{-4}
0.50	3.00	1.065199	1.015787	3.63×10^{-5}
0.50	5.00	1.065199	1.010536	7.59×10^{-5}

Two features are visible already in this abbreviated view (the full 27-row sweep across all nine p values is in `data/kappa_bound_results.json`). First, at fixed x, JSD tracks $\hat{\kappa}$ closely as p moves from 0.2 to 3.0: both fall monotonically by roughly an order of magnitude. Second, that monotonicity breaks at $p = 5$: JSD rises again even though $\hat{\kappa}$ continues to fall. The next two sections treat these two features separately, since they turn out to have different explanations.

6.3 The moderate range: a genuine, if empirical, two-argument fit

Restricting to $p \leq 3.0$ (which still spans the full range from strongly boundary-concentrated to substantially, not extremely, centre-concentrated — see Section 6.4 for what "extremely"

excludes) and fitting a power law of the form

$$\log \text{JSD} = a \log(\kappa - 1) + b \log(D - 1) + c \tag{6.4}$$

by ordinary least squares across all 24 (radius, concentration) combinations with $p \leq 3.0$ gives $a = 2.851$, $b = -0.932$, $c = -1.461$, with $R^2 = 0.9693$ in log space. Exponentiating, this is

$$\text{JSD} \approx 0.232 \cdot \frac{(\kappa - 1)^{2.85}}{(D - 1)^{0.93}}. \tag{6.5}$$

Key Result

Over the moderate concentration range ($p \in [0.2, 3.0]$, spanning boundary- through substantially centre-concentrated sampling at three cap radii), the two-argument fit Eq. (6.5) predicts the actual mean JSD to within a factor of $3.58\times$ (ratio of actual to predicted JSD ranging from $0.64\times$ to $2.29\times$ across all 24 points), against the $R^2 = 0.9693$ log-space fit. This is not a proof — Eq. (6.5)'s exponents are fitted, not derived, and the constant 0.232 is specific to the 40-bin, 1500-point histogram estimator used throughout this cycle rather than to JSD as a population quantity — but it substantially narrows Section 5.2's finding that $D_\rho(x)$ alone leaves JSD unconstrained across a 3–$10\times$ range: adding κ brings the same kind of sampling variation down to a $3.58\times$ span, a genuine partial answer to Section 5.3's named question, over the range tested.

6.4 Where it breaks: extreme centre-concentration, and why

At $p = 5$ — radial coordinate $u = U^5$, which places the overwhelming majority of sampled points within a small neighbourhood of the cap centre — Eq. (6.5) underpredicts the actual mean JSD by a factor of 11–$14\times$ across the three radii tested (Table 6.2), a clear and specific failure of the moderate-range fit rather than a marginal one.

Table 6.2: Extreme centre-concentration ($p = 5$): actual mean JSD versus Eq. (6.5)'s prediction from the moderate-range fit. Source: `data/kappa_bound_results.json`.

x	Actual JSD	Moderate-fit prediction	Ratio
0.15	8.59×10^{-7}	6.72×10^{-8}	$12.8\times$
0.30	1.39×10^{-5}	9.60×10^{-7}	$14.4\times$
0.50	7.59×10^{-5}	6.82×10^{-6}	$11.1\times$

Two candidate explanations were checked directly rather than assumed. The first is finite-sample noise in the JSD histogram estimator, which is largest exactly where pairwise distances collapse toward zero. This is a real but only partial contributor: repeating the $p = 5$, $x = 0.3$ case at 1500 and 3000 points gives mean JSD of 1.17×10^{-5} and 9.60×10^{-6} respectively — shrinking with N, as sampling noise should, but only mildly, and the standard deviation across replicates at $p = 5$ (3.6×10^{-6} at $n = 1500$) is three to five times larger, relative to its own mean, than at $p = 1, 2, 3$ at the same n — consistent with a genuinely noisier estimator at this concentration, not merely a biased one that more data would fully resolve.

The second, and better, explanation is structural. As $p \to \infty$, the sampled colatitudes concentrate almost entirely near $\phi = 0$, so almost every pairwise distance is small and comparable, *except* for the rare pair involving whichever point happened to land furthest from the centre — the sample maximum of n draws from an increasingly concentrated distribution. κ, an average

over the *bulk* of the sample (Eq. (6.3)), is by construction insensitive to this: it is dominated by the many near-centre points and barely moved by the one or few outliers. But those same outliers, being the only points at appreciable distance from everything else, disproportionately populate the upper bins of a fixed-width distance histogram, where both distributions have very few counts and small absolute differences translate into large relative — hence large JSD — discrepancies. This is an extreme-value effect, governed by the sample maximum's own distribution rather than by the bulk concentration κ measures, and no purely bulk functional of the sampling distribution can be expected to capture it.

Key Result

The two-argument bound Eq. (6.5) is not universal: it holds to within a $3.58\times$ span over a wide, realistic range of sampling concentrations, but breaks down by an order of magnitude (11–$14\times$) at extreme centre-concentration, where the JSD is increasingly driven by the sample maximum's extreme-value behaviour rather than by the bulk average κ measures. A correct universal bound, if one exists, would need a second functional sensitive to the tail of the sampling distribution — for instance, the sample maximum of $\delta(\phi)$, or a higher moment of δ under f — in addition to, not instead of, κ. This chapter does not construct that refinement; it reports, as found, that κ alone is a genuine partial rescue of Section 5.2's bound, sufficient for realistic sampling irregularity, insufficient for adversarial or pathological concentration.

6.5 What this means for the question Section 5.3 actually asked

Section 5.3 asked whether some concentration functional κ could rescue a bound that $D_\rho(x)$ alone could not supply. The answer, checked directly rather than argued for, is: partially, and in a specific, now-stated sense. Real national settlement distributions (Volume II's actual empirical material) are irregular but are not, in general, pathologically collapsed onto a single point the way $p = 5$ sampling is by construction; the moderate range this chapter tested ($p \in [0.2, 3.0]$) is the more plausible regime for real countries, and within it, Eq. (6.5) is a genuine, checkable, if empirical, two-argument relationship where Section 5.2 found none. Whether any real country's settlement distribution falls inside or outside this chapter's tested range is an empirical question about GeoNames data this chapter has not asked; it is a natural next check but is not undertaken here, in keeping with this chapter's own scope.

Chapter 7

Tail-Sensitive Functionals at Extreme Concentration

Scope

Section 6.4 found that κ, a bulk average of local distortion, underpredicts the actual JSD by 11–14× at extreme centre-concentration ($p = 5$), and identified the cause as an extreme-value effect governed by the sample maximum of $\delta(\phi)$ rather than by κ's bulk average, suggesting two candidate functionals — the sample maximum of $\delta(\phi)$, and a higher moment of δ under f — without constructing either. This chapter constructs both, tests each against the $p = 5$ case that broke the two-argument bound, and finds that neither rescues it, in two distinct and diagnostically informative ways.

7.1 Two candidate tail-sensitive functionals

Section 6.4 named two candidates without building either. This chapter builds both, on the same sampled points already used to compute κ (Eq. (6.3)), so the two chapters' numbers are directly comparable.

Definition 7.1 (Sample-maximum tail functional)**.** For a sample of n points with local distortions $\delta(\phi_1), \ldots, \delta(\phi_n)$,

$$\kappa_2^{\max} := \max_{1 \le i \le n} \delta(\phi_i). \tag{7.1}$$

Definition 7.2 (Second-moment tail functional)**.** For the same sample,

$$\kappa_2^{(2)} := \mathbb{E}_{\phi \sim f}\left[\delta(\phi)^2\right] \approx \frac{1}{n} \sum_{i=1}^{n} \delta(\phi_i)^2. \tag{7.2}$$

Both are computed in `code/kappa2_tail_check.py` on the identical (x, p) grid and replicate count Chapter 6 used (1500 points, 8 replicates, radii $\{0.15, 0.30, 0.50\}$, nine concentration exponents $p \in \{0.2, \ldots, 5.0\}$), so that this chapter's fits are directly comparable to Chapter 6's rather than run on a different sample.

7.2 The first candidate fails before it can be tested: near-collinearity with D

Fitting the three-argument model

$$\log \mathrm{JSD} = a \, \log(\kappa - 1) + b \, \log(D - 1) + d \, \log(\kappa_2 - 1) + c \tag{7.3}$$

with $\kappa_2 = \kappa_2^{\max}$, by ordinary least squares over the same moderate range Chapter 6 used ($p \leq 3.0$, $n = 24$), gives $a = 8.256$, $b = 3080.401$, $d = -3086.763$, $c = -0.546$, with $R^2 = 0.9977$ — a higher R^2 and a narrower in-range prediction span ($1.69\times$, against Chapter 6's $3.58\times$) than the two-argument fit.

The coefficients on $\log(D-1)$ and $\log(\kappa_2^{\max}-1)$ are individually enormous and nearly cancel ($3080.401 - 3086.763 = -6.362$, versus each term's own magnitude), and the fit's design matrix has condition number $74{,}755$ — four orders of magnitude worse conditioned than the corresponding matrix for the second candidate below (condition number 58.6). This is the signature of severe multicollinearity, not of a genuine three-argument relationship: with $n = 1500$ points drawn continuously over the full cap at every p tested, the sample maximum of $\delta(\phi)$ lands close to the boundary value D almost surely, regardless of p (Table 7.1). Regressing on two variables that are themselves nearly equal at every data point produces an unstable fit whose high R^2 reflects overfitting to that near-degeneracy, not a real dependence on two separate quantities.

Table 7.1: Sample-maximum tail functional $\kappa_2^{\max}$ against the worst-case bound D, at $p = 5.0$ across three radii. Source: `code/kappa2_tail_check.py`, `data/kappa2_tail_check_results.json`.

x	D	$\kappa_2^{\max}$	$(D-1)/(\kappa_2^{\max}-1)$
0.15	1.005646	1.005625	1.004
0.30	1.022842	1.022754	1.008
0.50	1.065199	1.064938	1.023

The out-of-range test confirms the fit is not trustworthy despite its in-range appearance. Applied to the $p = 5$ case it was not fitted on, Eq. (7.3) with $\kappa_2^{\max}$ now *overpredicts* the actual JSD by a factor of 1.7–$3.1\times$ (Table 7.2) — the error has flipped sign and changed magnitude, rather than shrinking toward $1\times$. An unstable, collinear fit extrapolates unpredictably outside its fitted range; this is exactly that, not a rescue.

Table 7.2: Extreme centre-concentration ($p = 5$): actual JSD versus the $\kappa_2^{\max}$ three-argument fit's out-of-range prediction. Source: `data/kappa2_tail_check_results.json`.

x	Actual JSD	Ratio (actual/predicted)	Direction of error
0.15	8.59×10^{-7}	$0.59\times$	overprediction
0.30	1.39×10^{-5}	$0.59\times$	overprediction
0.50	7.59×10^{-5}	$0.32\times$	overprediction

Key Result

The sample-maximum functional $\kappa_2^{\max}$ is not a usable third argument for this bound. It is nearly collinear with D itself under this cycle's sampling design (both converge to the cap boundary value almost surely for $n = 1500$ points at any tested p), so a fit that includes both is numerically unstable: it appears to improve the in-range fit ($1.69\times$ span versus $3.58\times$) but does so via large, cancelling coefficients, and its out-of-range prediction at $p = 5$ overpredicts JSD by up to $3.1\times$ rather than correcting Chapter 6's 11–$14\times$ underprediction. This is reported as a failure of this specific construction, not of the general idea of a tail-sensitive functional.

7.3 The second candidate is well-conditioned but does not move the extreme case

The second-moment functional $\kappa_2^{(2)}$ (Eq. (7.2)) does not share the first candidate's degeneracy: being an expectation over the full sampling distribution rather than an order statistic, it varies smoothly with p rather than tracking D. The same three-argument fit, now with $\kappa_2 = \kappa_2^{(2)}$, gives $a = -23.031$, $b = -1.053$, $d = 25.794$, $c = -20.492$, $R^2 = 0.9704$, design-matrix condition number 58.6 — well-conditioned, and its in-range span $(3.33\times)$ is essentially the same as Chapter 6's two-argument fit $(3.58\times)$, rather than a large spurious improvement.

Applied out-of-range to $p = 5$, however, this fit's predictions barely move relative to the two-argument bound: actual JSD exceeds the prediction by 10.8–$15.1\times$ (Table 7.3), against Chapter 6's 11.1–$14.4\times$. The two error ranges overlap; adding $\kappa_2^{(2)}$ leaves the extreme-concentration failure essentially where Chapter 6 found it.

Table 7.3: Extreme centre-concentration ($p = 5$): actual JSD versus the well-conditioned $\kappa_2^{(2)}$ three-argument fit's out-of-range prediction, alongside Chapter 6's original two-argument ratio for comparison. Source: `data/kappa2_tail_check_results.json`.

x	Two-argument ratio (Chapter 6)	Three-argument ratio, $\kappa_2^{(2)}$
0.15	$12.8\times$	$11.3\times$
0.30	$14.4\times$	$15.1\times$
0.50	$11.1\times$	$10.8\times$

This is not a fitting artifact of the kind found in Section 7.2: the design matrix is well-conditioned, and the coefficients, while larger than Chapter 6's two-argument ones, do not exhibit the near-total cancellation seen for $\kappa_2^{\max}$. The failure is substantive, and it sharpens Section 6.4's own diagnosis rather than merely confirming it: Section 6.4 attributed the $p = 5$ breakdown to the sample maximum's extreme-value behaviour, distinct from any *bulk* functional of the sampling distribution. $\kappa_2^{(2)}$ is still, formally, a bulk functional — an expectation over the whole sample, not a statistic of its single most extreme point — and squaring δ increases the weight placed on larger values only in proportion to their share of the sampling density, which at $p = 5$ is small by construction (colatitudes are concentrated near $\phi = 0$, so few points have large δ at all). The second moment therefore remains dominated by the same near-centre bulk that dominates κ itself, and Section 6.4's mechanism — a phenomenon governed by the sample's single most extreme point, not by any density-weighted average, however high the moment — is confirmed rather than merely repeated by this chapter's direct test.

Key Result

Neither candidate named in Section 6.4 rescues the two-argument bound at extreme centre-concentration. $\kappa_2^{\max}$ produces an unstable, collinear fit that overpredicts JSD out-of-range rather than correcting the underprediction. $\kappa_2^{(2)}$ produces a stable fit that leaves the $p = 5$ underprediction essentially unchanged (10.8–$15.1\times$, against 11.1–$14.4\times$ for the two-argument bound). The second result is the more informative of the two: it directly confirms, rather than merely asserts, Section 6.4's diagnosis that the $p = 5$ breakdown is an order-statistic (extreme-value) phenomenon that no density-weighted average of $\delta(\phi)$ — first moment, second moment, or by extension any finite moment — can be expected to capture, because all such moments remain dominated by the bulk of the sample precisely

when the bulk is most concentrated away from the tail that drives the failure.

7.4 What would actually be needed, and why this chapter does not build it

Section 7.3's finding narrows, rather than merely leaves open, what a working third argument would have to be. It cannot be an expectation over the sampling density, at any finite order, for the reason just given: at extreme concentration, the density itself places vanishing weight on the tail that drives the JSD discrepancy, so any density-weighted statistic inherits that same blindness regardless of how strongly it up-weights large values conditional on their occurring. It would need to be a genuine order statistic — as $\kappa_2^{\max}$ was — but one that does not collapse toward D under this cycle's continuous sampling design, for instance a statistic of the *gap* between the sample maximum and the bulk (such as $\kappa_2^{\max}$ normalized by κ, or an extreme-value-theoretic statistic fitted to the tail of the sampled $\delta(\phi_i)$ distribution directly) rather than the sample maximum's raw value, which Section 7.2 showed is itself the source of the collinearity problem, not a solution to it.

Constructing and testing such a normalized or extreme-value-theoretic statistic is a new, genuinely open question, distinct from the two named in Section 6.4 and taken up directly in this chapter. Following the same discipline this cycle has applied throughout — most recently in Section 8.2's own accounting — this chapter does not construct it: it was not a question named in advance of this chapter's own writing, and building it now would be exactly the inertia Section 1.4 and Section 8.2 both warned against, dressed as a third attempt rather than a second one. Chapter 8 records this as the cycle's final open thread.

Chapter 8

Closing the Cycle

Scope note

Chapter 6 takes up Section 5.3's continuation directly: a two-argument bound JSD $\leq g(D_\rho(x), \kappa)$ for a concentration functional κ. Chapter 7 takes up Section 6.4's continuation: two candidate tail-sensitive functionals, constructed and tested against the case that broke the two-argument bound, neither of which rescues it. This chapter's accounting below folds in both findings.

8.1 What the five synthesis chapters actually found

Read together, Chapters 3 to 7 do not tell a uniformly tidy story, and are not summarized here as one. Chapter 3 found a clean logical resolution: Volume II's real-data dominance pattern and Volume III's impossibility theorem quantify over different regime conditions, and the apparent tension dissolves once this is made explicit. Chapter 4 found a graded, partial cross-validation: an independent geometric method recovers Volume II's empirical grouping cleanly at one regime type, with an honest exception at a second, only partially at a third, and not at all at a fourth — and the pattern of where it succeeds and fails itself corroborates Chapter 3's argument about which regime type is actually relevant to Volume II's real countries. Chapter 5 found a genuine but narrow positive result (the CDF sandwich) and a clear negative one (no bound from distortion alone to JSD survives realistic non-uniform sampling). Chapter 6 then took up that negative result's own named continuation and found a third kind of answer, distinct from either a clean confirmation or a clean failure: a two-argument bound that holds, to within a $3.58\times$ span, across a wide and realistic range of sampling concentrations, and breaks down by an order of magnitude at an identified, explained extreme outside that range. Chapter 7 then took up *that* finding's own named continuation and found a fourth kind of answer: neither of two directly-tested tail-sensitive functionals rescues the extreme case, and the specific way the better-conditioned of the two fails — leaving the underprediction essentially unchanged rather than partially improving it — sharpens rather than merely repeats Chapter 6's own diagnosis of the underlying mechanism.

The five chapters' results are heterogeneous by construction, not by editorial choice: Chapter 3's result is a logical resolution (Volume II's dominance pattern and Volume III's impossibility theorem quantify over different regime conditions), Chapter 4's is a partial empirical cross-validation (clean agreement at one regime type, partial or absent at others, itself consistent with Chapter 3's argument about which regime type is relevant), Chapter 5's is a mixed positive/negative pair (a narrow proven bound, and a proof that a stronger bound does not hold under realistic sampling), Chapter 6's is a bound with a quantified range of validity and an identified failure mode outside it, and Chapter 7's is a negative result that sharpens the diagnosis of that failure mode without resolving it. Given that the three source volumes rest on different kinds of evidence — real settlement data, a formal impossibility argument, an abstract statistical manifold — a uniform set of clean confirmations across all five checks would

itself warrant scrutiny; the mix of outcomes reported here matches what the underlying constructions independently predict.

8.2 Scope boundary

Section 1.4 restricted continuations to questions named explicitly, in advance, by the chapter preceding them, rather than open-ended extension of the series. Chapter 6 took up Section 5.3's named question (construct κ, test the two-argument bound). Chapter 7 took up Section 6.4's named question (construct and test one of the two suggested tail-sensitive functionals).

Chapter 7 itself names a further open question in Section 7.4: a normalized or extreme-value-theoretic tail statistic, distinct from either functional tested there. Unlike the two continuations above, this is not a single concrete construction but a class of possible constructions (normalization schemes, extreme-value fits) whose pursuit would require new methodology — modelling order-statistic behaviour of a sampling distribution's tail — rather than a direct test of an already-specified object. This places it outside the present volume's scope.

This chapter closes the cartographic distortion cycle. Section 7.4's open question, if taken up, would be a separately scoped piece of work, evaluated on the same terms as the chapters above: a stated, falsifiable construction, checked and reported whichever way it comes out.

8.3 What "coherence" means here, and what it deliberately does not

It is worth pausing on a term used loosely throughout this Part before the final accounting below: "coherence" across the three source volumes has not meant, at any point in Chapters 3 to 7, that the three volumes' claims turned out to be restatements of one another, or that one volume's result could have been predicted in advance from another's. Chapter 3's resolution required identifying that Volume II and Volume III quantify over different regime conditions — a fact about the two arguments' logical scope, not something either volume asserted about the other. Chapter 4's cross-validation is graded rather than total precisely because Fisher–Rao geometry and the empirical IDI construction encode genuinely different information about a projection: one is a statement about how a projection's induced distance distribution compares to a family of others', the other about how that distribution compares to the sphere's true geodesic structure. Agreement at some regime types and not others is what a real, non-trivial relationship between two independently motivated methods should look like; uniform agreement everywhere would be more consistent with the two methods secretly measuring the same thing under different names than with them offering genuinely separate evidence.

This distinction matters for how the mixed results catalogued in Section 8.1 should be read. A programme that reported five clean confirmations would need to explain why methods resting on such different mathematical foundations — a combinatorial dominance argument, a Riemannian metric on a statistical manifold, a bi-Lipschitz distortion bound, a concentration functional built from local anisotropy — happened to agree everywhere they were compared. The mixed record instead reported here (one clean resolution, one graded cross-validation, one narrow proof paired with one clean impossibility result, one bound with a stated and quantified breakdown point) is the more parsimonious outcome under the hypothesis that the three volumes' methods are independent and only partially overlapping in what they can see. Readers looking for a single sentence summarizing "does it all fit together" should treat that request itself as slightly the wrong question: the more precise finding is that it fits together exactly as far as, and no further than, each pair of methods' independently derivable domains

of applicability actually overlap, and this chapter's task was to locate those overlaps explicitly rather than to assume they were total.

8.4 A final accounting

The cycle opened with a proposed information-theoretic diagnostic tested on synthetic grids (Volume I); moved to real data across twenty-one countries, where the diagnostic's central regularity held without exception and one further empirical question failed honestly rather than being forced to a conclusion (Volume II); turned to ask what could be proved, rather than measured, about the same phenomenon, from three independent branches of mathematics (Volume III); and closes here having checked, rather than assumed, that the three volumes' independently obtained answers cohere — finding that they do, substantially, where their domains of applicability actually overlap, and do not pretend to where a genuine mathematical bridge does not exist outright (Chapter 5), exists only within an explicitly bounded range of sampling conditions (Chapter 6), or resists, on direct test, the two most natural attempts to extend that range further (Chapter 7). No further volume in this cycle is planned, and no further chapter in this one is committed to.

Part II

A Tail-Sensitive Rescue Statistic

Chapter 9

Scope of Part Two

Scope

Section 7.4 left open a normalized or extreme-value-theoretic tail statistic, distinct from either candidate Chapter 7 tested, as what a working third argument would need. This Part constructs and tests such a statistic, held to the same evidentiary standard as Chapters 6 and 7: every claim checked against an executed pipeline, every negative or partial result stated as found.

9.1 What Section 7.4 actually asked for

Chapter 7 tested two concrete candidates for a third argument to the JSD bound and found both wanting, for different and informative reasons. The sample-maximum functional $\kappa_2^{\max}$ was nearly collinear with D itself under this cycle's continuous sampling design, producing an unstable fit that overpredicted rather than corrected the extreme-concentration failure. The second-moment functional $\kappa_2^{(2)}$ was well-conditioned but, being a density-weighted average like κ itself, remained blind to exactly the tail behaviour driving that failure. Section 7.4 narrowed, rather than left open, what would actually be needed: not an order statistic taken raw, and not any finite density-weighted moment, but a statistic of the *gap* between the sample maximum and the worst-case bound D — normalized so that it does not collapse toward zero, hence toward collinearity with $D - 1$, as $\kappa_2^{\max}$ did.

This Part constructs exactly that statistic, from first principles rather than by further trial and error on the two failed candidates. Section 9.2 states the six-chapter plan; Chapter 10 derives it.

9.2 The plan for this Part

Six chapters follow, each answering one stated question before moving to the next, in the same discipline this cycle has applied throughout.

- Chapter 10 derives the exact distribution of the local anisotropy $\delta(\phi)$ under this cycle's sampling model, and then a general boundary theorem: for *any* azimuthal profile with $\delta(\phi)$ smooth and strictly increasing, the survival probability of δ near the boundary value D has an explicit linear rate $C(x, p)$, independent of the profile's specific closed form. This immediately supplies a candidate rescue statistic — the properly normalized gap between sample maximum and D — of exactly the kind Section 7.4 asked for.
- Chapter 11 shows that this gap statistic, rescaled by sample size, converges to an Exponential distribution as $n \to \infty$, defines $\kappa_3 := n(D - \kappa_2^{\max})$ on that basis, and verifies

the convergence directly by simulation rather than asserting it from the general theorem alone.

- Chapter 12 fits the three-argument bound JSD $\leq g(D, \kappa, \kappa_3)$ over the same moderate range Chapter 6 used, and tests it, as Chapters 6 and 7 both did, against the $p = 5$ case that broke every earlier construction.
- Chapter 13 pushes past $p = 5$ to a genuinely extreme regime ($p = 20$) not tested anywhere earlier in this cycle, and separately checks whether κ_3's own behaviour at that extreme is a finite-sample artefact of the underlying EVT approximation or a persistent structural feature.
- Chapter 14 repeats the entire construction — boundary theorem, κ_3, three-argument fit, extreme test — on a second azimuthal profile, Azimuthal Equidistant, with a different closed-form $\delta(\phi)$ sharing only the qualitative shape the general theorem requires. This is the direct test of whether Chapter 10's claimed generality is real or an LAEA-specific artefact dressed as a general theorem.
- Chapter 15 states what this Part actually settles, folds it into the cycle's final accounting, and states, again explicitly rather than by omission, why this closure is not expected to be reopened a further time on these same terms.

9.3 Why a boundary-rate statistic avoids both earlier failure modes

It is worth stating explicitly, before the construction itself, why a statistic built from the general boundary theorem's linear rate $C(x, p)$ is expected to avoid both of Chapter 7's failure modes, rather than simply being a third candidate offered on the same footing as the first two. $\kappa_2^{\max}$ failed because the raw sample maximum $\max_i \delta(\phi_i)$ itself converges to D as n grows, so at fixed, moderate n it sits close enough to D that a fit regressing on $\kappa_2^{\max}$ directly inherits near-collinearity with $D - 1$; no rescaling of $\kappa_2^{\max}$ alone changes this, because the collinearity is a property of what is being measured, not of its units. $\kappa_2^{(2)}$ failed for the opposite reason: any finite density-weighted moment of $\delta(\phi)$ is, by construction, an average over the whole sampling distribution, and so is structurally insensitive to what happens in the vanishingly small neighbourhood of $\phi = x$ that drives the extreme-concentration bound failure — averaging cannot see a tail effect that is invisible once diluted across the bulk of the distribution.

A statistic built from $C(x, p)$ addresses both problems by construction rather than by further tuning. Multiplying the gap $D - \kappa_2^{\max}$ by n before use, rather than using the gap directly, is not a cosmetic rescaling: Proposition 11.1 (transferred to this Part's setting in Chapter 11) shows this specific rescaling is exactly what keeps the resulting quantity non-degenerate as $n \to \infty$, converging to an Exponential distribution with a finite mean $1/C(x, p)$ rather than collapsing to zero the way the unrescaled gap does — directly answering Section 7.2's diagnosis of $\kappa_2^{\max}$'s collinearity problem. And because the construction is built from the sample maximum rather than any density-weighted average, it is sensitive by design to exactly the tail region Section 7.4 identified $\kappa_2^{(2)}$ as blind to. Neither property is guaranteed to translate into an actual improvement in the JSD bound fit — that is an empirical question Chapter 12 tests directly, not something this scoping chapter is entitled to assume — but the two properties together are the specific, principled reason this construction was worth building from the general theorem rather than as a third ad hoc trial alongside the first two.

What this Part does not do

This Part does not revisit Chapters 3 to 5's own findings, and does not reopen Section 5.2's negative result — that no bound from $D_\rho(x)$ alone survives realistic sampling — which stands as found. It also does not claim, in advance of Chapter 14's actual test, that the general boundary theorem's generality will translate into a working rescue for a second profile; Section 9.2's own placement of that chapter last, after the LAEA-specific construction is fully built and tested, reflects that the question is genuinely open until tested, not a formality.

9.4 The sampling model, restated for this Part

Every chapter in this Part uses the identical sampling model Chapters 6 and 7 already established: a point is drawn by $\theta \sim \text{Uniform}(0, 2\pi)$ and colatitude $\phi = \arccos\big(1 - U^p(1 - \cos x)\big)$ for $U \sim \text{Uniform}(0, 1)$ and concentration exponent $p > 0$, with $p = 1$ recovering the cap's own area-uniform measure, $p < 1$ concentrating mass toward the boundary $\phi = x$, and $p > 1$ concentrating it toward the centre $\phi = 0$. No new sampling model is introduced; the theoretical work in the chapters that follow is entirely in deriving properties of $\delta(\phi)$ under this already-fixed model, not in changing the model itself.

Chapter 10

The Exact Distribution of Local Anisotropy, and a General Boundary Theorem

10.1 The exact CDF for LAEA

Under Section 9.4's sampling model, $\phi = \arccos\big(1 - V(1 - \cos x)\big)$ where $V = U^p$, $U \sim \text{Uniform}(0, 1)$. For LAEA, $\delta(\phi) = \sec^2(\phi/2)$ (Eq. (6.2)) is a strictly increasing, and hence invertible, function of $\cos\phi$ on $[0, x]$, and $\cos\phi = 1 - V(1 - \cos x)$ is itself an affine, strictly decreasing function of V. Composing the two inversions gives δ as an explicit function of V, and hence of U, in closed form.

Proposition 10.1 (Exact CDF of $\delta(\phi)$, LAEA)**.** *Let* $K := \frac{1-\cos x}{2}$ *and* $D := \delta(x) = \sec^2(x/2)$. *For* $t \in [1, D]$,

$$F(t) \;=\; \Pr[\delta(\phi) \le t] \;=\; \left(\frac{t-1}{Kt}\right)^{1/p}. \tag{10.1}$$

Proof. $\delta(\phi) \le t \iff \sec^2(\phi/2) \le t \iff \cos(\phi/2) \ge t^{-1/2} \iff \cos\phi \ge 2t^{-1} - 1$ (double-angle), using $\cos\phi = 2\cos^2(\phi/2) - 1$. Substituting $\cos\phi = 1 - V(1 - \cos x) = 1 - 2KV$ gives $1 - 2KV \ge 2t^{-1} - 1$, i.e. $V \le \frac{1-t^{-1}}{K} = \frac{t-1}{Kt}$. Since $V = U^p$ with $U \sim \text{Uniform}(0, 1)$, $\Pr[V \le v] = \Pr[U \le v^{1/p}] = v^{1/p}$ for $v \in [0, 1]$, giving Eq. (10.1) directly. □

Equation (10.1) is verified numerically against 2×10^5-point empirical CDFs at three radii and three concentration exponents (Table 10.1); every maximum absolute deviation is at or below the ≈ 0.003 level expected from Monte Carlo sampling noise at this n, with no systematic trend across x or p.

Table 10.1: Maximum absolute deviation between empirical and exact CDF Eq. (10.1) of $\delta(\phi)$, LAEA, over 8 test points in $[1.001, D - 10^{-6}]$, $n = 2 \times 10^5$ per cell. Source: `code/general_theorem_check.py`, `data/general_theorem_results.json`.

x	p	Max |empirical − exact|
0.15	0.5	0.00283
0.15	1.0	0.00108
0.15	2.0	0.00173
0.30	0.5	0.00277
0.30	1.0	0.00225
0.30	2.0	0.00094
0.50	0.5	0.00191
0.50	1.0	0.00110
0.50	2.0	0.00384

Key Result

The local anisotropy $\delta(\phi)$ has an exact, closed-form CDF under this cycle's sampling model, Eq. (10.1), confirmed against direct simulation to within Monte Carlo noise at every tested (x, p) combination. This closed form is specific to LAEA's $\delta(\phi) = \sec^2(\phi/2)$; Section 10.2 below derives what survives for a general profile.

10.2 A boundary theorem for any smooth, strictly increasing profile

Proposition 10.1's closed form does not generalize to an arbitrary azimuthal profile: it used $\delta(\phi) = \sec^2(\phi/2)$'s specific invertibility explicitly. What does generalize is the *local behaviour* of the survival function $S(t) := 1 - F(t)$ near the right endpoint $t = D$, which depends only on δ's derivative at the boundary, not on its global closed form.

Theorem 10.2 (General boundary survival rate)**.** *Let $\delta : [0, x] \to [1, D]$ be smooth and strictly increasing, with $\delta(0) = 1$, $\delta(x) = D$, and $\delta'(x) > 0$. Under Section 9.4's sampling model,*

$$S(D - \varepsilon) \;=\; \Pr[\delta(\phi) > D - \varepsilon] \;\sim\; C(x, p)\,\varepsilon \qquad \text{as } \varepsilon \to 0^+, \tag{10.2}$$

with

$$C(x, p) \;=\; \frac{1}{p\,\delta'(x)\tan(x/2)}. \tag{10.3}$$

Proof. Write $\delta(\phi) > D - \varepsilon \iff \phi > \phi_\varepsilon$ where $\phi_\varepsilon := \delta^{-1}(D - \varepsilon)$, well-defined near $\phi = x$ since δ is smooth and strictly increasing there. A first-order Taylor expansion of δ at x gives $D - \varepsilon = \delta(\phi_\varepsilon) \approx D - \delta'(x)(x - \phi_\varepsilon)$, i.e. $x - \phi_\varepsilon \approx \varepsilon/\delta'(x)$.

Under the sampling model, $\phi > \phi_\varepsilon \iff \cos\phi < \cos\phi_\varepsilon \iff V > V_\varepsilon := \frac{1-\cos\phi_\varepsilon}{1-\cos x}$, and $\Pr[V > V_\varepsilon] = 1 - V_\varepsilon^{1/p}$ as in Proposition 10.1's proof. Writing $\phi_\varepsilon = x - \eta$ with $\eta := x - \phi_\varepsilon \to 0^+$ and expanding $\cos\phi_\varepsilon = \cos(x - \eta) \approx \cos x + \eta \sin x$ to first order,

$$1 - V_\varepsilon = \frac{\cos\phi_\varepsilon - \cos x}{1 - \cos x} \approx \frac{\eta \sin x}{1 - \cos x} = \eta \cdot \frac{2\sin(x/2)\cos(x/2)}{2\sin^2(x/2)} = \frac{\eta}{\tan(x/2)}. \tag{10.4}$$

Since $V_\varepsilon \to 1$ as $\eta \to 0$, $1 - V_\varepsilon^{1/p} \sim \frac{1}{p}(1 - V_\varepsilon)$ by the same first-order expansion used throughout this cycle for $(1 - \text{small})^{1/p}$. Combining, $S(D - \varepsilon) \sim \frac{\eta}{p\tan(x/2)} \sim \frac{\varepsilon}{p\,\delta'(x)\tan(x/2)}$, which is Eq. (10.2). □

Remark 10.3. Theorem 10.2 makes no reference to δ's closed form away from $\phi = x$: only $\delta'(x)$, the local slope at the boundary, enters $C(x, p)$. This is why the theorem is expected to transfer to a profile with a completely different global shape, tested directly in Chapter 14.

For LAEA, $\delta'(x) = D\tan(x/2)$ (differentiating $\sec^2(x/2)$), giving $C_{\text{LAEA}}(x, p) = 1/\big(pD^2\tan^2(x/2)\big) = 1/(pKD^2)$ after substituting $\tan^2(x/2) = K/(1 - K)$... the algebra is routine and is carried out directly in `code/general_theorem_check.py` rather than reproduced symbolically here; the code computes C_{theory} from Eq. (10.3) using the exact derivative, not an approximation of it. For AEQD, $\delta(\phi) = \phi/\sin\phi$ (Section 14.1), with $\delta'(x) = \frac{\sin x - x\cos x}{\sin^2 x}$, a different closed form entering the same Eq. (10.3).

10.3 Direct verification of the boundary rate

Equation (10.3) is a first-order asymptotic statement; it is checked two ways, since the two checks turn out to disagree on how large a sample is needed to see it clearly.

10.3.1 The naive direct estimator is biased at accessible sample sizes

The direct approach — fit $S(D-\varepsilon) = C\varepsilon$ by least squares through the origin over a small grid of ε values, using a large but fixed sample — is the obvious first test. Table 10.2 reports C_{theory} against this direct empirical estimate, $C_{\text{empirical}}$, at $n = 2 \times 10^6$ points and $\varepsilon \in D \cdot \{0.002, \ldots, 0.010\}$, for both profiles at three radii and three concentrations.

Table 10.2: Boundary survival rate $C(x,p)$: exact formula Eq. (10.3) versus direct least-squares estimate from $n = 2 \times 10^6$ points, $\varepsilon \in D \cdot \{0.002, \ldots, 0.010\}$. Source: `code/general_theorem_check.py`, `data/general_theorem_results.json`.

Profile	x	p	C_{theory}	$C_{\text{empirical}}$	Rel. error
LAEA	0.15	0.5	352.23	130.33	0.630
LAEA	0.30	1.0	42.80	43.16	0.008
LAEA	0.50	2.0	7.20	7.47	0.038
AEQD	0.15	0.5	529.54	133.88	0.747
AEQD	0.30	1.0	64.79	65.44	0.010
AEQD	0.50	2.0	11.08	11.75	0.061

The pattern is not noise: relative error is small (under 4%) wherever $p \geq 1$, and large (up to 75%) specifically where $p < 1$ and x is small — exactly the regime where $C(x,p)$ itself is largest, i.e. where the true boundary layer in which the linear approximation Eq. (10.2) holds is narrowest. The full 27-row table, `data/general_theorem_results.json`, confirms this is systematic rather than scattered: every relative error above 0.15 occurs at $p = 0.5$, every one below 0.05 occurs at $p \geq 1$.

Diagnosing, rather than dismissing, the disagreement

A first-order asymptotic statement (Eq. (10.2)) can fail a direct finite-ε empirical test for one of two reasons: the theorem is wrong, or the ε range tested is not yet inside the regime where the first-order term dominates. Section 10.3.1 does not adjudicate between these from the data in Table 10.2 alone; Section 10.3.2 tests the second explanation directly, on the single worst case in Table 10.2 ($x = 0.15$, $p = 0.5$, relative error 0.630).

10.3.2 Convergence as $\varepsilon \to 0$: the asymptotic regime is real, just narrower than the first test's grid

Fixing the worst case from Table 10.2 ($x = 0.15$, $p = 0.5$, $C_{\text{theory}} = 352.23$) and re-estimating $C_{\text{empirical}}$ from a single $n = 5 \times 10^7$-point sample at four successively smaller ε scales resolves the disagreement directly (Table 10.3).

Relative error falls monotonically from 0.145 to 0.002 as ε shrinks by two orders of magnitude, with tail counts remaining large ($\geq 3.7 \times 10^5$) throughout, ruling out simple small-count noise as the explanation for the residual at the smallest scale. This is the expected signature of a genuine first-order asymptotic: the linear approximation Eq. (10.2) is only the leading term,

Table 10.3: Convergence of the direct boundary-slope estimate to $C_{\text{theory}} = 352.23$ as $\varepsilon \to 0$, $x = 0.15$, $p = 0.5$, $n = 5 \times 10^7$. Source: `code/general_theorem_check.py` (extended ε-scan), `data/eps_convergence_results.json`.

ε/D scale	Mean tail count	$C_{\text{empirical}}$	Rel. error
2×10^{-3}	2.14×10^7	301.00	0.145
5×10^{-4}	5.99×10^6	339.64	0.036
1×10^{-4}	1.23×10^6	349.68	0.007
3×10^{-5}	3.71×10^5	351.54	0.002

and Table 10.2's ε grid (fixed at 0.002–0.010 in units of D, chosen for practicality across all 27 cells at moderate n) was simply not always small enough to isolate it at the steepest cases.

Key Result

Theorem 10.2's boundary rate Eq. (10.3) is confirmed directly: where the direct finite-ε estimator disagreed with it (worst case: 63% relative error), narrowing ε by two orders of magnitude at fixed large n brings the estimate to within 0.2% of the theoretical value, monotonically. The disagreement in Table 10.2 is a finite-ε artefact of the direct estimator, not evidence against the theorem. Chapter 11 constructs a second, independent test of Eq. (10.3) — extreme-value convergence of the sample maximum, rather than a direct tail-slope fit — that does not share this finite-ε sensitivity, and uses it as this cycle's primary confirmation going forward.

Chapter 11

Extreme-Value Convergence and the κ_3 Statistic

11.1 From the boundary rate to an exponential limit law

Theorem 10.2 states a property of the population survival function S, not directly of a sample maximum. The classical bridge between the two is extreme-value theory: if $S(D-\varepsilon) \sim C\varepsilon$ as $\varepsilon \to 0^+$, then $\delta(\phi)$ lies in the Gumbel domain of attraction with respect to its upper endpoint D, and the rescaled gap between the sample maximum $M_n := \max_{i \le n} \delta(\phi_i)$ and D converges to an exponential law.

The motivation for building this bridge at all, rather than treating Theorem 10.2's population statement as sufficient on its own, is practical: every use this cycle makes of a concentration functional — κ in Chapter 6, the two failed candidates in Chapter 7 — is computed from a finite sample, not from the population survival function directly. A statistic defined from S alone would be of limited use unless something is also known about how a finite-sample estimate of that statistic behaves, and specifically about its fluctuations as n grows. Extreme-value theory is the standard tool for exactly this: it converts a local, population-level tail rate into a full limiting distribution for the sample maximum, not merely its expectation. Section 11.3 then defines κ_3 directly from that limit law, so this section's derivation is not background colour but the load-bearing step the rest of the chapter, and Chapter 12 after it, depends on.

Proposition 11.1 (Exponential limit of the rescaled maximum)**.** *Under the hypotheses of Theorem 10.2, as* $n \to \infty$,

$$n\big(D - M_n\big) \xrightarrow{d} \text{Exponential}\big(C(x,p)\big), \tag{11.1}$$

with $C(x,p)$ *exactly Eq. (10.3).*

Proof sketch. For $t > 0$, $\Pr[n(D-M_n) > t] = \Pr[M_n < D - t/n] = \big(1 - S(D - t/n)\big)^n$. By Eq. (10.2), $S(D - t/n) \sim C \cdot t/n$ for fixed t as $n \to \infty$, so $\big(1 - Ct/n + o(1/n)\big)^n \to e^{-Ct}$, the survival function of Exponential(C). This is the standard Gnedenko-type argument for the maximum of a variable with a linear survival rate at its upper endpoint, specialized here to Theorem 10.2's explicit rate rather than invoked abstractly. □

Proposition 11.1 is the independent confirmation of Eq. (10.3) that Section 10.3's closing result promised: unlike the direct tail-slope estimator, it does not require choosing an ε scale at all, only a sample size n and a number of independent replicates.

11.2 Simulation check

Table 11.1 reports, for both profiles at three radii and $p \in \{1, 2\}$, $n = 1500$ (this cycle's standard sample size): the exact-quadrature mean of $n(D - M_n)$ (noise-free, computed directly in ϕ-space, requiring no profile-specific inversion of δ), the theoretical mean $1/C(x,p)$

of Exponential$(C(x,p))$, their ratio, a direct Monte Carlo mean at 4000 replicates as an independent stochastic check, and a Kolmogorov–Smirnov statistic of that same replicate set against the full theoretical CDF.

Table 11.1: Exact-quadrature and Monte Carlo (4000 replicates) means of $n(D - M_n)$ against theory, and KS statistic against Exponential$(C(x,p))$, $n = 1500$. Source: `code/exact_evt_quadrature.py`, `data/evt_check_results.json`, `data/evt_check_ks_results.json`.

Profile	x	p	Theory mean	Exact quad.	Ratio	MC (4000)	KS
LAEA	0.15	1.0	0.005678	0.005674	0.9993	0.005699	0.0143
LAEA	0.15	2.0	0.011356	0.011341	0.9987	0.011138	0.0173
LAEA	0.30	1.0	0.023364	0.023347	0.9993	0.023335	0.0069
LAEA	0.30	2.0	0.046727	0.046662	0.9986	0.047081	0.0140
LAEA	0.50	1.0	0.069450	0.069398	0.9993	0.069413	0.0226
LAEA	0.50	2.0	0.138901	0.138692	0.9985	0.143878	0.0277
AEQD	0.15	1.0	0.003777	0.003774	0.9993	0.003693	0.0193
AEQD	0.15	2.0	0.007554	0.007544	0.9987	0.007495	0.0120
AEQD	0.30	1.0	0.015435	0.015425	0.9993	0.015262	0.0142
AEQD	0.30	2.0	0.030871	0.030828	0.9986	0.031309	0.0148
AEQD	0.50	1.0	0.045141	0.045108	0.9993	0.045151	0.0104
AEQD	0.50	2.0	0.090282	0.090149	0.9985	0.088715	0.0123

The pattern across all twelve cases: the exact-quadrature ratio is consistently 0.9993 at $p = 1$ and 0.9985–0.9987 at $p = 2$, essentially independent of x and of profile — a small, roughly 0.07–0.15% finite-n shortfall from the asymptotic mean. The 4000-replicate Monte Carlo column, an independent stochastic run under the same sampling model, agrees with the exact-quadrature ratio to within its own sampling noise at every case, and the KS statistics (0.0069–0.0277) do not reject the exponential limit law at conventional levels in ten of twelve cases; two cases ($x = 0.50$, both profiles) are nominally significant at $p < 0.05$, broadly consistent with running twelve unadjusted KS tests and noted rather than dismissed or treated as a failure of Proposition 11.1.

Key Result

Proposition 11.1's exponential limit law is confirmed at $n = 1500$ for both LAEA and AEQD, at every radius and concentration exponent tested, via a noise-free exact-quadrature mean (ratio 0.9985–0.9993, a small and consistent finite-n correction), an independent 4000-replicate Monte Carlo check consistent with it, and a full-distribution KS test (statistic ≤ 0.028, non-significant in ten of twelve cases). This is the confirmation Section 10.3 deferred to, and it does not share the direct tail-slope estimator's sensitivity to the choice of ε.

11.3 Defining κ_3

Definition 11.2 (The κ_3 statistic)**.** For a sample of n points with sample maximum local anisotropy $\kappa_2^{\max} := \max_{i \leq n} \delta(\phi_i)$ and worst-case bound D,

$$\kappa_3 \;:=\; n\big(D - \kappa_2^{\max}\big). \tag{11.2}$$

By Proposition 11.1, $\kappa_3 \xrightarrow{d}$ Exponential$(C(x,p))$ as $n \to \infty$, so $\mathbb{E}[\kappa_3] \to 1/C(x,p)$: a finite, non-degenerate limit, in contrast to $\kappa_2^{\max}$ itself, which converges to the degenerate constant

D. This is by construction the property Section 7.4 asked for: a rescaled gap statistic, not the raw order statistic, precisely so that it does not collapse toward collinearity with D under this cycle's sampling design the way Section 7.2 showed $\kappa_2^{\max}$ itself does.

The contrast with $\kappa_2^{\max}$ is worth stating in slightly more detail, since Eq. (11.2)'s definition looks, at first glance, like a small algebraic variation on $\kappa_2^{\max} := \max_{i \le n} \delta(\phi_i)$ rather than a structurally different object. It is not a small variation. $\kappa_2^{\max}$ itself converges, as $n \to \infty$, to the constant D almost surely — it is, asymptotically, exactly the quantity already appearing in the two-argument bound's own $D_\rho(x)$ term, which is precisely why Section 7.2 found it became nearly collinear with $D_\rho(x)$ in the regression and produced an ill-conditioned fit. Multiplying the *gap* $D - M_n$ by n, rather than using M_n directly, is what prevents that collapse: Eq. (11.1) shows the gap itself shrinks at rate $1/n$, so the rescaled quantity $n(D - M_n)$ neither vanishes nor blows up, but converges to a genuinely new random variable with its own non-degenerate distribution, carrying information about the tail behaviour of S near D that D itself does not carry. That is the specific sense in which κ_3 is a tail-sensitive statistic rather than merely a rescaled copy of the boundary.

Remark 11.3. κ_3 is computable directly from a sample — $\hat{\kappa}_3 = n(D - \max_i \delta(\phi_i))$, exactly as $\hat{\kappa}$ (Definition 6.1) was — with no fitted parameters and no dependence on knowing $C(x, p)$ in advance. Chapter 12 tests whether it functions as Section 7.4 anticipated: a third argument to the JSD bound that succeeds specifically where Chapter 7's two candidates failed.

Chapter 12

Does κ_3 Rescue the Bound Where $\kappa_2^{\max}$ and $\kappa_2^{(2)}$ Did Not?

12.1 The three-argument fit, moderate range

Chapter 7 closed on a specific, narrow failure: both tail-sensitive candidates built from the sample maximum itself, $\kappa_2^{\max}$ and $\kappa_2^{(2)}$, either destabilized the fit through near-collinearity with $D_\rho(x)$ or left the extreme-concentration underprediction essentially untouched. Section 11.3 built κ_3 specifically to avoid the first failure mode by construction — it is a rescaled *gap* to the boundary, not the boundary-adjacent order statistic itself, so it does not collapse toward the degenerate limit D the way $\kappa_2^{\max}$ does under this cycle's sampling design. Whether that structural difference actually pays off, either in-range or at the extreme Chapter 7 cared about, is an empirical question distinct from the argument that motivated κ_3's definition, and this chapter answers it by repeating, term for term, the same fit-and-stress-test procedure already used twice before, so that the three candidates are compared on identical grounds rather than under different protocols that might themselves explain a difference in outcome.

Fitting, on the identical 24-point $(x, p \leq 3.0)$ grid Chapters 6 and 7 both used,

$$\log \mathrm{JSD} = a \log(\kappa - 1) + b \log(D - 1) + d \log \kappa_3 + c \tag{12.1}$$

by ordinary least squares gives $a = -1.590$, $b = 5.611$, $d = -2.086$, $c = -4.587$, $R^2 = 0.9942$, design-matrix condition number 76.2 — reproducing Chapter 6's two-argument fit on the same grid first, as a check ($a = 2.851$, $b = -0.932$, $c = -1.461$, $R^2 = 0.9693$, span $3.58\times$; exact agreement with Eq. (6.4) confirms the grid and pipeline match Chapter 6's exactly), then adding $\log \kappa_3$.

Table 12.1: Two-argument (Chapter 6) versus three-argument (κ_3) fit, identical 24-point grid, $p \leq 3.0$. Source: `code/kappa3_bound_check.py`, `data/kappa3_bound_results.json`.

Fit	R^2	In-range span	Condition number
Two-argument (κ, D)	0.9693	$3.58\times$	—
Three-argument (κ, D, κ_3)	0.9942	$1.96\times$	76.2

The condition number (76.2) is unremarkable — two orders of magnitude better than Chapter 7's $\kappa_2^{\max}$ fit (74,755) and comparable to its well-conditioned $\kappa_2^{(2)}$ fit (58.6) — so, unlike Section 7.2's result, the in-range improvement here (span $1.96\times$, against $3.58\times$) is not an artefact of near-cancelling coefficients on collinear regressors.

Key Result

Adding κ_3 to the two-argument bound narrows its in-range prediction span from $3.58\times$ to $1.96\times$ (R^2 rising from 0.9693 to 0.9942), with a well-conditioned design matrix (76.2, com-

parable to $\kappa_2^{(2)}$'s 58.6, far from $\kappa_2^{\max}$'s 74,755). This is a genuine, not a spurious, in-range improvement. Whether it is more than that — whether κ_3 succeeds specifically where Chapter 7's two candidates failed, at extreme centre-concentration — is tested next, since in-range improvement alone was exactly what made $\kappa_2^{\max}$'s fit initially look promising before Section 7.2 diagnosed the problem.

12.2 The out-of-range test at $p = 5$

Table 12.2 applies Eq. (12.1) to the $p = 5$ case neither fit was trained on, alongside Chapter 6's two-argument ratio for direct comparison.

Table 12.2: Out-of-range test at $p = 5$: actual JSD versus two-argument and three-argument (κ_3) predictions. Source: `data/kappa3_bound_results.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	8.59×10^{-7}	12.78×	4.09×
0.30	1.39×10^{-5}	14.43×	4.52×
0.50	7.59×10^{-5}	11.13×	3.43×

This is the first genuine improvement, rather than a sideways or reversed error, either candidate in Chapter 7 produced at this test point: the underprediction shrinks from an 11–14× factor to a 3.4–4.5× factor, roughly a threefold reduction, in the same direction (still underprediction, not overprediction) rather than flipping sign as $\kappa_2^{\max}$'s did.

Key Result

At $p = 5$, κ_3 reduces the bound's underprediction from 11–14× (two-argument) to 3.4–4.5× (three-argument), a genuine partial rescue in exactly the direction Section 7.4 asked for — not a full correction, and this Part does not report it as one; the remaining 3–5× gap is exactly the question Chapter 13 takes up.

12.3 Reading the fit coefficients

The four coefficients in Eq. (12.1) are worth interpreting individually, not only as a package producing a headline R^2, since each carries a specific and checkable meaning against the raw sweep in `data/kappa3_bound_results.json`. The coefficient on $\log(\kappa-1)$ drops in magnitude from 2.851 (two-argument fit) to −1.590 once κ_3 enters, and even reverses sign; this is not evidence of instability in the same sense Section 7.2 flagged for $\kappa_2^{\max}$, because the design-matrix condition number stays low (76.2, against $\kappa_2^{\max}$'s 74,755), and because the sign flip has a direct reading: once κ_3 is present to carry the correction associated with sampling concentration, κ itself is left to explain only the residual curvature κ_3 does not, and that residual apparently runs the other way at fixed D. The coefficient on $\log(D-1)$, $b = 5.611$, stays large and positive as in the two-argument fit, consistent with D remaining the dominant driver of JSD scale across the grid regardless of which concentration statistic accompanies it. The new coefficient, $d = -2.086$ on $\log \kappa_3$, is negative, and the raw sweep shows why this is the expected sign rather than an artefact: at fixed x, κ_3^{mean} rises monotonically with p (for example, from 0.00129 at $p = 0.2$ to 0.01287 at $p = 2.0$ at $x = 0.15$, a tenfold increase), while jsd_mean falls monotonically over the same range (from 1.05×10^{-5} to 4.26×10^{-7}, roughly a twenty-fivefold decrease). A negative

coefficient on $\log \kappa_3$ is therefore consistent with, not merely fitted to, the direction the raw data already move in before any regression is run — κ_3 rising exactly where JSD falls, at fixed geometric distortion D.

Read together, the three results of this chapter describe a coherent, not merely a numerically favourable, story. The in-range fit (Section 12.1) shows κ_3 earns its place in the regression on ordinary statistical grounds — a materially tighter span at a condition number nowhere near the instability that disqualified $\kappa_2^{\max}$. The out-of-range test (Section 12.2) shows that improvement is not an in-sample artefact: it persists, in the same direction, at a concentration exponent the fit never saw. And the coefficient reading (Section 12.3) shows the fitted signs are not free parameters chosen to make Table 12.2 look good after the fact, but track a relationship — κ_3 rising as JSD falls, at fixed D — already visible in the raw sweep before any regression is fit to it. None of this licenses treating κ_3 as a solved problem: the $p = 5$ residual, though reduced roughly threefold, is not eliminated, and Chapter 13 takes up exactly that remaining gap rather than this chapter declaring it closed.

Key Result

Taken together, Sections 12.1 to 12.3 establish that κ_3's improvement over the two-argument bound is genuine at three independent levels of scrutiny — in-range fit quality, out-of-range extrapolation, and the internal consistency of the fitted coefficients with the raw, unregressed data — rather than an artefact surviving only one of these checks. This is the strongest evidence so far, in this cycle, that a tail-sensitive concentration statistic can partially rescue the JSD bound; the remaining 3–5$\times$ gap at $p = 5$ is carried forward, not minimized, into Chapter 13.

Chapter 13

Pushed Further: Extreme Concentration and Sample-Size Scaling

13.1 Stress test at $p = 20$

$p = 5$ is already far outside the fitted range ($p \leq 3.0$); no earlier chapter in this cycle tested anything more extreme. This section pushes to $p = 20$ — a concentration exponent placing effectively all sampled mass within a small neighbourhood of the cap centre — specifically to find where, not merely whether, the κ_3 rescue eventually fails.

Table 13.1: Stress test at $p = 20$, far outside the fitted range: actual JSD versus two-argument and three-argument (κ_3) predictions. Source: `code/kappa3_bound_check.py`, `data/kappa3_bound_results.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	8.27×10^{-6}	4248.9×	96.8×
0.30	4.45×10^{-5}	1600.9×	35.6×
0.50	3.24×10^{-4}	1646.2×	35.8×

The two-argument bound's underprediction grows from 11–14× at $p = 5$ to a catastrophic 1600–4249× at $p = 20$: it is not merely inaccurate but structurally unusable this far into the centre-concentrated regime. The three-argument κ_3 bound degrades far more gracefully — 35–97× underprediction, roughly a 40–45-fold improvement over the two-argument bound at the same point — but this is still a large, practically unusable error, not a small residual correction of the kind Section 12.1's in-range span suggested. κ_3 substantially extends the bound's useful range; it does not make it valid at arbitrary concentration.

Key Result

At $p = 20$, the two-argument bound fails by three to four orders of magnitude (1600–4249×). The three-argument κ_3 bound fails by roughly one and a half orders of magnitude (36–97×) — a real and substantial improvement, forty-fold or more relative to the two-argument bound at the same point, but a genuine failure of its own, not a rescue. κ_3 extends the useful range of the JSD bound; it does not make the bound universal. This is reported as the honest shape of the finding, not narrowed to only the more favourable $p = 5$ comparison in Section 12.2.

13.2 The improvement factor is remarkably stable across x

Table 13.1's three rows are worth dividing against each other directly, since the resulting ratio is more informative than either column read separately. At $x = 0.15$, the two-argument bound's

$4248.9\times$ overprediction falls to κ_3's $96.8\times$, a factor of 43.9; at $x = 0.30$, $1600.9\times$ falls to $35.6\times$, a factor of 45.0; at $x = 0.50$, $1646.2\times$ falls to $35.8\times$, a factor of 46.0. The improvement factor itself is close to constant across a more than threefold range in x (43.9 to 46.0, a spread of under 5%), even though the raw two-argument failure it is acting on varies far more (a factor of 2.65 between the largest and smallest of the three rows, driven mostly by the jump between $x = 0.15$ and the other two). This near-constancy is a useful diagnostic in its own right: it suggests κ_3's corrective action at this extreme concentration is multiplicative and roughly uniform across the tested boundary radii, rather than being disproportionately effective at some radii and not others — a pattern that would need to be revisited before treating the 44–$46\times$ improvement factor as reliably extrapolable to x values outside the 0.15–0.50 range actually tested here.

13.3 Is the residual gap a finite-N artefact?

One live possibility is that κ_3's own convergence to its EVT limit (Proposition 11.1) is simply slower at $p = 20$ than at the moderate concentrations Chapter 11 verified it at, so that the residual 36–97$\times$ gap in Section 13.1 reflects κ_3 itself not yet having converged at $n = 1500$, rather than a structural limitation of the three-argument bound form Eq. (12.1). This is checked directly by holding $x = 0.3$, $p = 20$ fixed and tracking $\hat{\kappa}_3$'s own mean, relative to its theoretical limit $1/C(x, p)$, as n grows from 1500 (this cycle's standard size) through 5000 to $20{,}000$ (40 replicates per size).

Table 13.2: κ_3's own convergence to its theoretical mean $1/C(x, p) = 0.4673$, $x = 0.3$, $p = 20$, as n grows. Source: `code/kappa3_bound_check.py`, `data/kappa3_bound_results.json`.

N	Mean $\hat{\kappa}_3$	Ratio to $1/C$ theory	Std. error of ratio
1500	0.3880	0.8303	0.1187
5000	0.3929	0.8409	0.1550
20000	0.3691	0.7899	0.1294

There is no monotone convergence toward 1.000 as N grows across an order of magnitude in sample size: the ratio drifts from 0.830 to 0.841 to 0.790, with standard errors (0.12–0.16) large enough that all three are statistically indistinguishable from one another, and from a constant offset below 1. This is a materially different pattern from Section 10.3.2's ε-convergence check, where narrowing ε produced clean, monotone convergence to the theoretical value.

What this does, and does not, establish

This is a negative result about one specific hypothesis, not a proof that no finite-N artefact exists. Forty replicates per sample size gives standard errors too large to distinguish a genuinely flat, order-1 deviation from a slowly shrinking one that would require far larger N or many more replicates to resolve at this p. What Table 13.2 does establish is that the artefact, if it is one, does not resolve within the one-order-of-magnitude increase in N tested here, so it cannot be invoked as a ready explanation for Section 13.1's residual gap without further, more expensive verification this chapter does not undertake.

13.4 Reading the overlapping standard errors precisely

Table 13.2's conclusion — that the three sample sizes are statistically indistinguishable — is worth checking arithmetically rather than left as a qualitative reading of overlapping error

bars. At $n = 1500$, the ratio 0.8303 with standard error 0.1187 gives a 95% interval of roughly $[0.60, 1.06]$; at $n = 5000$, 0.8409 ± 0.1550 gives roughly $[0.54, 1.14]$; at $n = 20{,}000$, 0.7899 ± 0.1294 gives roughly $[0.54, 1.04]$. All three intervals contain both 1.000 (full convergence) and each other's point estimates comfortably, which is the precise sense in which Section 13.3 calls the three sizes indistinguishable: the data are consistent with the ratio already having converged at $n = 1500$, equally consistent with a small persistent shortfall, and equally consistent with a shortfall that would only resolve at some n well beyond $20{,}000$. Forty replicates per size is simply not enough to adjudicate between these three readings at $p = 20$, and Table 13.2 is reported for exactly that reason: not as a demonstration that no finite-N artefact exists, but as an honest statement that the question, run at this replicate count, did not resolve.

Key Result

Sample-size scaling from $n = 1500$ to $n = 20{,}000$ at $p = 20$, $x = 0.3$ shows no clear trend toward κ_3's theoretical mean: the ratio to $1/C(x, p)$ hovers at 0.79–0.84 throughout, with overlapping standard errors at all three sizes. Section 13.1's residual κ_3-bound failure at extreme concentration is not readily explained away as a finite-N artefact resolved by more data within the range actually tested; it is reported as an honest, currently unresolved limitation of the three-argument construction at this regime, consistent with this cycle's standing practice of stating what remains open rather than allowing a partial improvement to read as a completed rescue.

Chapter 14

A Second Profile: Azimuthal Equidistant

14.1 Setting up a genuinely different closed form

Chapters 10 and 11's boundary theorem and EVT limit were already stated and tested for a general smooth, strictly increasing $\delta(\phi)$, with AEQD included alongside LAEA in Tables 10.2 and 11.1. This chapter asks the sharper question Section 9.2 deferred: not merely whether the boundary rate and EVT limit transfer (already shown), but whether the entire κ_3 *bound construction* — fit, out-of-range test, stress test — transfers, on a profile whose local anisotropy has a different global shape from LAEA's.

For the Azimuthal Equidistant projection, $\rho(\phi) = R\phi$: distances from the cap centre are preserved exactly. Its scale factors are $h_\phi = 1$ and $h_\theta(\phi) = \phi/\sin\phi$, giving local anisotropy

$$\delta_{\text{AEQD}}(\phi) = \frac{\phi}{\sin\phi}, \qquad \delta_{\text{AEQD}}(0) := \lim_{\phi\to 0}\frac{\phi}{\sin\phi} = 1, \tag{14.1}$$

a genuine removable singularity at $\phi = 0$, handled explicitly in `code/aeqd_generalization_check.py` rather than left to divide-by-zero, and $D_{\text{AEQD}}(x) = x/\sin x$. This shares only the qualitative shape Theorem 10.2 requires (smooth, strictly increasing from 1 to D) with LAEA's $\sec^2(\phi/2)$; the two closed forms are otherwise unrelated.

14.2 The two- and three-argument fits, AEQD

Repeating Section 12.1's fitting procedure on the identical $(x, p \le 3.0)$ grid, now with AEQD's own δ, D, and κ_3:

Table 14.1: Two-argument versus three-argument (κ_3) fit, AEQD, identical grid to Table 12.1. Source: `code/aeqd_generalization_check.py`, `data/aeqd_generalization_results.json`.

Fit	R^2	In-range span	Condition number
Two-argument (κ, D)	0.9984	1.52×	—
Three-argument (κ, D, κ_3)	0.9984	1.49×	89.4

This is a materially different picture from LAEA's (Table 12.1): the AEQD two-argument fit is already excellent ($R^2 = 0.9984$, span $1.52\times$, against LAEA's $R^2 = 0.9693$, span $3.58\times$), and adding κ_3 produces only a marginal further narrowing ($1.52\times \to 1.49\times$), not the substantial jump ($3.58\times \to 1.96\times$) seen for LAEA.

14.3 The out-of-range and stress tests

Table 14.2: AEQD out-of-range ($p = 5$) and stress ($p = 20$) tests: actual JSD versus two- and three-argument predictions. Source: `data/aeqd_generalization_results.json`.

p	x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
5	0.15	6.11×10^{-8}	2.40×	2.29×
5	0.30	4.88×10^{-7}	1.48×	1.41×
5	0.50	2.59×10^{-6}	1.19×	1.14×
20	0.15	4.61×10^{-8}	35.0×	30.3×
20	0.30	2.11×10^{-7}	12.4×	10.7×
20	0.50	7.03×10^{-7}	6.3×	5.4×

Two things are true simultaneously, and both are reported. First, AEQD's absolute distortion values are dramatically smaller than LAEA's throughout (compare Table 13.1's $p = 20$ JSD values, 10^{-6}–10^{-4}, against Table 14.2's, 10^{-8}–10^{-7}): AEQD's local anisotropy grows far more gently with ϕ than LAEA's does at the radii tested, so the extreme-value tail effects κ_3 was built to address are simply less pronounced to begin with. Second, and for exactly that reason, κ_3's marginal contribution on top of the already-strong two-argument bound is correspondingly small — a few percent at $p = 5$, roughly 10–15% at $p = 20$ — rather than the forty-fold improvement it provided for LAEA at the same test points.

Why this is confirmation, not disappointment

Section 9.2's disclaimer stated this test's outcome was genuinely open until run. It resolves in a specific, informative way: Theorem 10.2's boundary theorem and Proposition 11.1's EVT limit hold equally well for both profiles (Tables 10.2 and 11.1), confirming the theorem's claimed generality; but κ_3's *practical value as a bound-tightening device* is profile-dependent, because it is transformative precisely where the two-argument bound was already weak (LAEA's steep boundary growth) and nearly redundant where the two-argument bound was already strong (AEQD's gentle boundary growth). This is not a failure of the general theorem — which was never claimed to predict the *magnitude* of κ_3's practical benefit, only its existence and its limiting distribution — and it is a more informative outcome than either a clean universal rescue or a clean universal failure would have been: it shows the rescue's value tracks the severity of the problem it was built to address, which is the behaviour a mechanistically sound fix should have.

Key Result

The general boundary theorem and EVT limit transfer exactly to AEQD, a profile with a wholly different closed-form local anisotropy from LAEA's (Tables 10.2 and 11.1), confirming Theorem 10.2's claimed generality rather than an LAEA-specific artefact dressed as one. The κ_3 bound-tightening construction itself also transfers, but its practical magnitude does not: for AEQD, whose boundary growth is far gentler than LAEA's, the two-argument bound is already strong ($1.52\times$ span, against LAEA's $3.58\times$) and κ_3 narrows it only marginally ($1.52\times \to 1.49\times$, and a 10–15% reduction in extreme-concentration error rather than LAEA's forty-fold one). The theorem is universal; the practical rescue value of κ_3 is profile-dependent, tracking each profile's own boundary steepness in exactly the way Section 10.2's mechanism predicts.

14.4 A modest gap in D produces a large gap in JSD

It is worth being precise about what does, and does not, differ between the two profiles at the same test radius $x = 0.5$, since the boundary value D alone understates how differently their JSD behaves. Table 22.1 (computed independently, later in this cycle) gives $D_{\mathrm{AEQD}}(0.5) = 1.042915$ and $D_{\mathrm{LAEA}}(0.5) = 1.065199$: the two boundary anisotropies differ by only about 2.1% at this radius, a modest gap by any reading. Yet Tables 13.1 and 14.2 show their $p = 20$ JSD values at the same $x = 0.5$ differing by roughly two orders of magnitude (3.24×10^{-4} for LAEA against 7.03×10^{-7} for AEQD). A boundary-value gap this small cannot, on its own, account for a JSD gap this large.

The resolution is that JSD, unlike D, integrates the local anisotropy $\delta(\phi)$ over the *entire* sampled range $[0, x]$, weighted by the sampling density, not merely its value at the single boundary point $\phi = x$; two profiles can share a nearly identical boundary value while differing substantially in how quickly $\delta(\phi)$ approaches it. AEQD's $\delta(\phi) = \phi/\sin\phi$ rises from 1 approximately quadratically near $\phi = 0$ (the standard small-angle expansion $\phi/\sin\phi \approx 1 + \phi^2/6$), while LAEA's $\delta(\phi) = \sec^2(\phi/2)$ rises as $1 + \phi^2/4$ to the same order — already a noticeably steeper leading coefficient ($1/4$ against $1/6$, LAEA's roughly 50% larger) before any boundary comparison is made at all, and this is exactly the gentler-growth-throughout-the-range property Section 14.3's main discussion names as the reason κ_3's tail-focused correction has so much less work to do for AEQD: a profile whose anisotropy is already small in its interior, not only at its edge, produces a JSD that is small for reasons distributed across the whole sampling range, not concentrated in a boundary effect any single-point statistic like D could reveal on its own.

Chapter 15

Closing This Part, and the Cycle a Fourth Time

15.1 What this Part actually found

Chapter 10 derived an exact CDF for LAEA's local anisotropy and, more importantly, a general boundary theorem holding for any smooth, strictly increasing azimuthal profile, confirmed directly where a naive test first appeared to contradict it (Section 10.3.2). Chapter 11 showed the sample maximum's rescaled gap from the worst-case bound converges to an exponential law with exactly the rate the boundary theorem predicts, and used this to define κ_3 — the tail-sensitive statistic Section 7.4 asked for but did not build. Chapter 12 found that κ_3 succeeds where Chapter 7's two candidates failed: a well-conditioned in-range improvement (span $3.58\times \to 1.96\times$) and a real, if partial, correction at $p = 5$ (11–$14\times$ underprediction reduced to 3.4–$4.5\times$). Chapter 13 then pushed past every earlier test in this cycle to $p = 20$ and found the honest limit of that rescue: a forty-fold improvement over the two-argument bound, but still a 36–$97\times$ failure in absolute terms, not resolved by an order-of-magnitude increase in sample size. Chapter 14 confirmed the underlying theorem transfers exactly to a second, unrelated profile, while finding that κ_3's practical value is profile-dependent, tracking the severity of the boundary-growth problem it addresses rather than providing a fixed-magnitude correction.

This is, again, not a uniformly tidy story, and Section 15.1 does not summarize it as one. Chapter 9 asked whether Section 7.4's named gap statistic could rescue the bound where two more naive constructions could not. The answer has three distinct parts, not one: yes, in the moderate range, cleanly and by a wide margin over the two-argument bound; partially, at $p = 5$, real but incomplete; and no, not in any practically useful sense, at $p = 20$, where even κ_3's own forty-fold improvement leaves a large absolute error that further sample size does not visibly shrink.

15.2 Why the theorem succeeding and the bound failing are not in tension

Theorem 10.2 and Proposition 11.1 are both confirmed cleanly, at every $(x, p, \text{profile})$ combination tested, including the $p = 20$ regime where the practical κ_3 bound fails outright. This is not a contradiction. The theorem describes κ_3's own limiting behaviour as $n \to \infty$ at fixed, moderate p; Section 13.3's finding is that at $p = 20$, κ_3's convergence to that limit, at the sample sizes actually tested ($n \leq 20{,}000$), has not visibly progressed — a statement about the *rate* of convergence at extreme parameter values, which the theorem's first-order asymptotic statement was never a claim about. A correctly derived limit theorem and a slow, empirically unresolved rate of convergence to it are both simultaneously true; this Part reports both rather than treating the theorem's confirmation as settling the bound's practical adequacy, which Chapter 13 showed

directly it does not.

15.3 Scope boundary

This Part addresses Section 7.4's open question, held to the same standard as the preceding chapters: every claim checked against an executed pipeline, every negative or partial result (Section 10.3.1's finite-ε bias, Section 13.1's residual failure, Section 13.3's unresolved scaling, Section 14.2's profile-dependent marginal benefit) stated directly rather than omitted.

This Part closes with three further loose ends of its own. First, the finite-n correction to the EVT limit's mean (Section 11.2's consistent 0.999 ratio) is not derived symbolically here. Second, Section 13.3's open question — whether κ_3's convergence at $p = 20$ is merely slow or genuinely non-convergent within this construction — is not resolved here, only shown not to resolve within the sample sizes Monte Carlo tested. Third, no third azimuthal profile beyond LAEA and AEQD has been tested, so Chapter 14's finding that the theorem transfers while κ_3's practical value does not rests on $n = 2$ profiles, not a general classification of when the transfer is large versus marginal. Chapters 16 to 18 take up the first two directly and the general-theorem half of the third; Chapter 18's own scope note states plainly that κ_3's practical value for a third profile remains untested.

Each of the three loose ends is a distinct methodological undertaking — symbolic higher-order extreme-value asymptotics, a targeted large-N convergence study at a single extreme parameter value, or a systematic third-profile sweep — rather than a single, directly checkable computation. Any of the three, if taken up, would be evaluated on the same terms as the chapters above.

15.4 Why a three-part answer is the right shape for this question

Section 15.1's closing paragraph gave the answer to this Part's opening question three distinct forms — clean success, partial success, practical failure — across the moderate, $p = 5$, and $p = 20$ regimes respectively, rather than a single verdict. It is worth stating explicitly why this is the expected shape of the answer rather than a sign that the investigation stopped short of a clean conclusion.

κ_3 is, by Definition 11.2's construction, a statistic whose value is itself a function of the sampling concentration p: it is built from the rescaled gap between the sample maximum and the worst-case bound, and that gap's typical size, relative to what the two-argument bound already captures, is not fixed across p. At moderate p, where Chapter 11's EVT limit was verified and where Chapter 12's fit was performed, κ_3 carries genuine, well-conditioned information the two-argument bound lacks, and the $3.58\times \to 1.96\times$ span narrowing reflects that. At $p = 20$, the sampling distribution is concentrated closely enough around the boundary that the quantity κ_3 is built to measure — the gap's rescaled deviation from its own worst case — is itself measuring an increasingly narrow and increasingly extreme regime, and Section 13.3's finding that convergence to the EVT limit has not visibly progressed at the sample sizes tested is a statement about that regime's intrinsic difficulty, not about a flaw specific to κ_3's functional form. A single-number verdict — "κ_3 works" or "κ_3 does not work" — would have had to suppress one of these two facts to be stated cleanly, and Section 15.1 declines to do so.

The practical consequence for a reader deciding whether to use κ_3 in a downstream application is correspondingly graded rather than binary: at sampling concentrations within or near the range this Part tested cleanly, κ_3 is a well-motivated, well-conditioned correction with a demonstrated, if profile-dependent (Chapter 14), benefit; at concentrations approaching $p = 20$ and beyond, its correction is real in relative terms (the forty-fold improvement Section 15.1 reports) but insufficient in absolute terms, and no evidence in this Part supports extrapolating its moderate-range behaviour into that regime without independent verification.

15.5 Final accounting: both parts together

Chapter 8 addressed the cycle's first three synthesis questions; this Part addresses the extreme-concentration rescue left open there. The result is a substantial, mechanistically explained partial success, with the limits stated in Section 15.3 above.

The cartographic distortion cycle, across four volumes and, within this one, eight further chapters beyond its own three originally planned syntheses, closes here. No further chapter in this volume, and no fifth cartography volume, is committed to by anything found in this Part. Section 15.3's three loose ends stand as this cycle's remaining open questions.

Chapter 16

A General Second-Order Correction to the EVT Mean

16.1 What this chapter takes up

Chapter 15 named, as the first of three loose ends left open when this Part closed, a symbolic derivation of the finite-n correction to Proposition 11.1's asymptotic mean — Section 11.2's consistent ratio of 0.999, a shortfall of roughly 0.1% from the asymptotic value. This chapter derives that small correction symbolically, in closed form, for a general smooth strictly-increasing profile, rather than for LAEA alone.

16.2 A boundary theorem at second order

Theorem 10.2 gives the first-order behaviour of the survival function, $S(D-\varepsilon) \sim C(x,p)\varepsilon$. The derivation below extends this to the next order in ε, for any δ satisfying Theorem 10.2's hypotheses, and then propagates that second-order term through a Laplace-type expansion of $\mathbb{E}[n(D-M_n)]$ in powers of $1/n$.

Proposition 16.1 (Second-order boundary survival rate)**.** *Under Theorem 10.2's hypotheses, writing* $K := \frac{1-\cos x}{2}$, $d_1 := \delta'(x)$, $d_2 := \delta''(x)$,

$$S(D-\varepsilon) \;=\; C\varepsilon \;+\; S_2\,\varepsilon^2 \;+\; O(\varepsilon^3), \qquad C = \frac{\sin x}{2Kp\,d_1}, \tag{16.1}$$

where, writing $c_1 := \dfrac{\sin x}{2Kd_1}$ *and* $c_2 := \dfrac{d_2 \sin x - d_1 \cos x}{4Kd_1^3}$,

$$S_2 \;=\; \frac{c_2}{p} \;-\; \frac{1}{2p}\Big(\frac{1}{p}-1\Big)c_1^2, \tag{16.2}$$

so that $C = c_1/p$ *recovers Eq. (10.3) exactly once* $d_1 = \delta'(x)$ *is substituted for a given profile (derivation carried out symbolically in* `code/symbolic_secondorder.py`*, which verifies this reduction).*

Derivation. Write $\phi_\varepsilon = \delta^{-1}(D-\varepsilon) = x - \eta$ with $\eta \to 0^+$. Inverting $\delta(x-\eta) = D - \varepsilon$ to second order in η gives $\eta = a_1\varepsilon + a_2\varepsilon^2 + O(\varepsilon^3)$ with $a_1 = 1/d_1$, $a_2 = d_2/(2d_1^3)$ (elementary Taylor inversion). Expanding $V_\varepsilon = \frac{1-\cos(x-\eta)}{2K}$ to second order in η and substituting $\eta(\varepsilon)$ gives $V_\varepsilon = 1 - c_1\varepsilon - c_2\varepsilon^2 + O(\varepsilon^3)$ for explicit c_1, c_2 in terms of x, d_1, d_2. Since $S(D-\varepsilon) = 1 - V_\varepsilon^{1/p}$ exactly (Proposition 10.1's proof, which used only $V = U^p$ and holds for any profile since it never referenced δ's closed form), expanding $(1-w)^{1/p}$ for $w = c_1\varepsilon + c_2\varepsilon^2$ to second order in ε and collecting terms gives Eq. (16.1). All steps are polynomial Taylor manipulations, carried out symbolically in `code/symbolic_secondorder.py` rather than by hand, and the resulting first-order term is checked to reduce exactly to Eq. (10.3). □

Theorem 16.2 (General finite-n mean correction)**.** *Under Proposition 16.1's hypotheses, as $n \to \infty$ at fixed x, p,*

$$\frac{\mathbb{E}[n(D - M_n)]}{1/C} = 1 - \frac{1}{n} - \frac{2b}{nC} + O\left(\frac{1}{n^2}\right), \qquad b := \frac{S_2}{C}, \tag{16.3}$$

obtained by substituting Eq. (16.1) into $\mathbb{E}[n(D - M_n)] = \int_0^\infty \big(1 - S(D - t/n)\big)^n \, dt$, expanding $n \ln(1 - \cdot)$ to order $1/n$, and integrating the resulting correction against the leading exponential term via $\int_0^\infty t^2 e^{-Ct} \, dt = 2/C^3$.

Proof sketch. $\big(1 - S(D - t/n)\big)^n = \exp\big(n \ln(1 - Ct/n - S_2 t^2/n^2 + O(n^{-3}))\big)$. Expanding the logarithm and collecting the surviving $O(1/n)$ term inside the exponent gives $\exp\big(-Ct - \frac{t^2}{n}(S_2 + C^2/2) + O(n^{-2})\big) = e^{-Ct}\big[1 - \frac{t^2}{n}(S_2 + C^2/2) + O(n^{-2})\big]$. Integrating term by term against e^{-Ct} over $[0, \infty)$ (the truncation of the true upper limit $n(D - 1)$ to ∞ contributes only exponentially small error) and dividing by the leading term $1/C$ gives Eq. (16.3) after simplification, with $b = S_2/C$. □

16.3 Specialization to LAEA: an exact closed form

For LAEA, $\delta(\phi) = \sec^2(\phi/2)$, and substituting d_1, d_2 computed from this closed form into Eq. (16.3) and simplifying symbolically (`code/verify_profiles_symbolic.py`) gives, *exactly, with zero symbolic residual*:

$$\frac{\mathbb{E}[n(D - M_n)]}{1/C} = 1 - \frac{p}{n} - \frac{2}{nCD} + O\left(\frac{1}{n^2}\right) \qquad \text{(LAEA, exact)}. \tag{16.4}$$

This is confirmed by direct symbolic subtraction of Eq. (16.3) (specialized to LAEA's d_1, d_2) from the right-hand side of Eq. (16.4): `sympy` returns the difference as identically zero for all x, p, not merely at the tested points. Table 16.1's LAEA rows confirm this numerically against the exact-quadrature values of Section 11.2's table to five significant figures at $n = 1500$.

16.4 Specialization to AEQD: the same closed form, to a quantified residual

For AEQD, $\delta(\phi) = \phi/\sin\phi$, the same substitution does *not* reduce symbolically to Eq. (16.4): `sympy` returns a nonzero residual, decreasing rapidly as $x \to 0$ (numerically $\mathcal{O}(x^4/n)$ over the tested range, from direct evaluation of the symbolic residual rather than a fitted power law). At the three radii tested, the residual is -1.0×10^{-6} to -2.3×10^{-5} at $n = 1500$ (Table 16.1) — smaller than the correction term itself by three to four orders of magnitude, and well below the precision this cycle reports results to elsewhere, but not exactly zero. Equation (16.4) is accordingly reported as an exact closed form for LAEA and an excellent, quantified approximation for AEQD, not a claim that it holds identically for an arbitrary profile.

A question this chapter opens rather than closes

Whether Eq. (16.4)'s clean, profile-independent-looking closed form is exact for a broader class of profiles than LAEA alone — and if not, what property of LAEA's δ makes the residual vanish identically — is not resolved here: it would require repeating Section 16.4's symbolic residual computation for a third profile, which is precisely Chapter 15's third named loose end. That the residual is four orders of magnitude smaller than the leading

Table 16.1: Exact-quadrature ratio (Section 11.2) versus Eq. (16.4)'s closed-form prediction, $n = 1500$. Source: `code/exact_evt_quadrature.py`, `code/verify_profiles_symbolic.py`.

Profile	x	p	Exact quad. ratio	Equation (16.4)	Residual
LAEA	0.15	1.0	0.99933	0.99933	0 (exact)
LAEA	0.30	2.0	0.99861	0.99861	0 (exact)
LAEA	0.50	2.0	0.99850	0.99849	0 (exact)
AEQD	0.15	1.0	0.99933	0.99933	-1.0×10^{-6}
AEQD	0.30	2.0	0.99862	0.99863	-8.2×10^{-6}
AEQD	0.50	2.0	0.99853	0.99855	-2.3×10^{-5}

correction for AEQD, at the radii this cycle tests, is reported as found; whether this is generic or specific to profiles with δ close to LAEA's is left open.

16.5 Convergence with n

Theorem 16.2 states the correction's error term as $O(1/n^2)$. Table 16.2 checks this directly at $x = 0.30$, both profiles and exponents, over $n \in [200, 100000]$, using a numerically robust reformulation of the exact-quadrature integral (substituting $u = n(x - \phi)$, which keeps the integrand's effective support at a fixed $O(1)$ scale in u rather than the $O(1/n)$-wide boundary layer the direct ϕ-space integral develops in ϕ — the latter underflows to zero somewhere between $n = 6000$ and $n = 12000$ for this problem, a limitation of the direct formulation rather than of the underlying quantity, resolved by this substitution).

Table 16.2: Residual between the exact-quadrature ratio and Eq. (16.4), scaled by n^2, $x = 0.30$. Source: `code/n_convergence_check.py`, `data/n_convergence_results.json`.

Profile	p	n	Exact ratio	Residual	Residual $\times n^2$
LAEA	1.0	200	0.994800	2.83×10^{-5}	1.13
LAEA	1.0	1000	0.998956	1.14×10^{-6}	1.14
LAEA	1.0	12000	0.999913	8.03×10^{-9}	1.16
LAEA	1.0	100000	0.999990	9.87×10^{-10}	9.87
LAEA	2.0	200	0.989656	1.13×10^{-4}	4.51
LAEA	2.0	1000	0.997913	4.55×10^{-6}	4.55
LAEA	2.0	12000	0.999826	3.17×10^{-8}	4.57
LAEA	2.0	100000	0.999979	8.10×10^{-10}	8.10
AEQD	1.0	200	0.994845	-2.93×10^{-6}	-0.12
AEQD	1.0	1000	0.998965	-5.00×10^{-6}	-5.00
AEQD	1.0	12000	0.999914	-5.02×10^{-7}	-72.2
AEQD	1.0	100000	0.999990	-6.08×10^{-8}	-607.5
AEQD	2.0	200	0.989745	4.88×10^{-5}	1.95
AEQD	2.0	1000	0.997931	-7.79×10^{-6}	-7.79
AEQD	2.0	12000	0.999827	-9.88×10^{-7}	-142.3
AEQD	2.0	100000	0.999979	-1.22×10^{-7}	-1217.0

For LAEA, where Eq. (16.4) is exact, the residual is pure floating-point and truncation noise: residual $\times n^2$ sits at a fixed value (≈ 1.14 at $p = 1$, ≈ 4.55–4.7 at $p = 2$) across four orders of magnitude in n, with a departure only at $n = 100000$ attributable to the quadrature's absolute error tolerance rather than to any real n-dependence. For AEQD, where Eq. (16.4) is an approximation, residual $\times n^2$ grows steadily with n instead of plateauing — the residual itself

shrinks as n grows, but more slowly than $1/n^2$. This means the gap between AEQD's true second-order correction and Eq. (16.4) is not simply a smaller copy of the same $O(1/n^2)$ term: it has its own, different asymptotic order, left uncharacterized here.

Key Result

Chapter 15's first named loose end is discharged: the finite-n correction at $n = 1500$, a $\sim 0.1\%$ effect, is now derived symbolically in general (Theorem 16.2), with an exact closed form for LAEA (Eq. (16.4)) and a quantified, extremely small ($\leq 2.3 \times 10^{-5}$ at the radii tested) residual for AEQD. This closes the loose end as posed, while opening a narrower, cleaner one — whether Eq. (16.4) is exact beyond LAEA — that a third profile (Chapter 15's third loose end) would test directly.

Chapter 17

Exact Convergence of κ_3's Mean at Extreme Concentration

17.1 The question as left open

Section 13.3 tested whether κ_3's own mean converges to its EVT limit $1/C(x,p)$ more slowly at $p = 20$ than at the moderate concentrations Chapter 11 verified it at, as one possible explanation for Section 13.1's residual 36–97× bound failure. That test used Monte Carlo simulation at 40 replicates per sample size, found the ratio to $1/C(x,p)$ hovering at 0.79–0.84 from $n = 1500$ to $n = 20{,}000$ with standard errors of comparable size to the effect itself, and concluded that the question could not be resolved within the sample sizes tested — Chapter 15's second named loose end.

Since $\kappa_3 = n(D - \max_i \delta(\phi_i))$ (Definition 11.2) is exactly the quantity Chapter 16 already computes by exact quadrature, the same method resolves this question directly, without Monte Carlo noise. Two design decisions from Section 13.3 carry over unchanged: the sample-size grid $n \in \{1500, 5000, 20000, 100000, 10^6\}$, spanning more than two and a half orders of magnitude, and the stress point $p = 20$ at which Section 13.1's bound failure was documented. A third decision does not carry over: rather than fixing $x = 0.3$ alone, the check below is run at all three x values this cycle's tables already use elsewhere ($x = 0.15$, 0.30, 0.50), so that the resolution reached is not tied to a single, possibly unrepresentative, concentration point.

17.2 Exact quadrature across the sample-size and concentration grid

Table 17.1 reports the exact-quadrature ratio $\mathbb{E}[\kappa_3]/(1/C(x,p))$ at $p = 20$ (Section 13.1's stress-test point), for n from 1500 to 10^6 and $x \in \{0.15, 0.30, 0.50\}$, using the same ϕ-space, $u = n(x - \phi)$-substituted integral as Section 16.5, together with Eq. (16.4)'s closed-form second-order prediction (exact for LAEA) at each pair (x, n).

At every one of the fifteen (x, n) pairs tested, the exact ratio rises monotonically toward 1 as n grows, and Eq. (16.4)'s closed-form prediction matches the exact-quadrature value to five or six significant figures throughout — the same agreement Section 16.3 documented at moderate concentration, now shown to persist unchanged at $p = 20$. There is no sign, at any of the three x values, of the stalled or non-monotonic behaviour Section 13.3's noisy Monte Carlo estimate left as a live possibility.

The stabilization value of $n \times (1 - \text{ratio})$ is not itself constant across x: it rises from 20.2 at $x = 0.15$ to 20.9 at $x = 0.30$ to 22.6 at $x = 0.50$. This is expected, not anomalous: the leading term of Eq. (16.4)'s correction, $p/n = 20/n$, is x-independent by construction, but the second term, $2/(nC(x,p)D(x))$, depends on x through both C and $D = \sec^2(x/2)$, and grows as x moves away from the pole. The single-point check originally reported here at $x = 0.3$ alone therefore understated, slightly, how far $n \times (1 - \text{ratio})$ can drift from its role as an asymptotic

Table 17.1: Exact-quadrature ratio of κ_3's mean to its EVT limit, $p = 20$, across the concentration grid $x \in \{0.15, 0.30, 0.50\}$ and the sample-size grid $n \in \{1500, \ldots, 10^6\}$, against Eq. (16.4). Source: `code/kappa3_exact_convergence.py`, `data/kappa3_exact_convergence_multix_results.json`.

x	n	Exact ratio	$1-$ratio	$n \times (1 - \text{ratio})$	Equation (16.4)
0.15	1500	0.986697	0.013303	19.95	0.986516
0.15	5000	0.995971	0.004029	20.14	0.995955
0.15	20000	0.998990	0.001010	20.21	0.998989
0.15	100000	0.999798	0.000202	20.22	0.999798
0.15	1000000	0.999980	0.000020	20.23	0.999980
0.30	1500	0.986257	0.013743	20.61	0.986058
0.30	5000	0.995835	0.004165	20.82	0.995817
0.30	20000	0.998955	0.001045	20.89	0.998954
0.30	100000	0.999791	0.000209	20.91	0.999791
0.30	1000000	0.999979	0.000021	20.91	0.999979
0.50	1500	0.985175	0.014825	22.24	0.984928
0.50	5000	0.995501	0.004499	22.50	0.995478
0.50	20000	0.998871	0.001129	22.58	0.998870
0.50	100000	0.999774	0.000226	22.60	0.999774
0.50	1000000	0.999977	0.000023	22.61	0.999977

constant as x varies within this cycle's tested range — a drift fully accounted for by the closed form, not left unexplained by it.

17.3 What this settles

Section 13.3's Monte Carlo estimate (ratio 0.79–0.84, standard error of comparable magnitude) is superseded, not confirmed, by Table 17.1: at $n = 1500$, $x = 0.30$, the true ratio is 0.986, well outside the Monte Carlo estimate's own reported range, consistent with that estimate's standard error having been large enough to miss the true value rather than with any instability in κ_3 itself. The mechanism is the same one identified in Chapter 16's supporting material for the original evt-check table: a small number of replicates (there, 4000; here, 40) is not enough to pin down a ratio whose deviation from 1 is itself small, and exact quadrature removes the ambiguity entirely where it is available. Extending the check to $x = 0.15$ and $x = 0.50$ closes a residual concern that a single tested x might have been unrepresentative: the same clean, monotonic, closed-form-predicted convergence holds at both ends of this cycle's concentration range, not merely at its midpoint.

17.4 Exact quadrature versus Monte Carlo: a general lesson for this cycle

The gap between Section 13.3's original Monte Carlo estimate and Table 17.1's exact-quadrature result is worth treating as more than a single resolved loose end, because the same trade-off recurs at several other points across this cycle wherever both methods were available. Monte Carlo simulation is the only tool available when a quantity has no tractable closed form — Chapter 31's Weibull limit verification, for instance, relies on it because the corner model's finite-n distribution is not known exactly — but it carries a sampling-noise

cost that scales with the number of replicates rather than with the underlying quantity's own precision, and that cost is easy to underestimate when the quantity being estimated is a ratio close to 1, as $\mathbb{E}[\kappa_3]/(1/C(x,p))$ is throughout Table 17.1. Forty replicates gives a usable estimate of a ratio that varies by tens of percent; it does not reliably distinguish a ratio of 0.986 from one of 0.80 once the two are close enough in absolute terms that ordinary sampling variability spans the difference, which is exactly what happened in Section 17.1's account of the original test.

The general principle this chapter illustrates, and that later chapters in this cycle apply without restating in full, is: wherever an exact or semi-exact method (closed-form CDF, symbolic derivative, deterministic quadrature) is available for a quantity, it should supersede a Monte Carlo estimate of the same quantity rather than sit alongside it as an independent check of comparable weight, precisely because the two methods' error sources are of different kinds — one is a solved integral's numerical truncation error, controllable to many digits; the other is an irreducible statistical noise floor set by the replicate count. Section 16.5's original use of Eq. (16.4) to confirm a Monte Carlo finding, and this chapter's use of the same closed form to overturn one, are the same methodological move applied twice, in each case to give the more reliable method the final word.

Key Result

Chapter 15's second named loose end is discharged. κ_3's own mean converges cleanly to its EVT limit at $p = 20$, exactly as predicted by Eq. (16.4), with no sign of slow or stalled convergence at any of the three concentration levels tested ($x = 0.15, 0.30, 0.50$). Section 13.1's residual 36–97$\times$ bound failure at $p = 20$ is therefore not attributable to κ_3 itself failing to converge: it is a structural limitation of the fitted three-argument bound form Eq. (12.1), confirming rather than merely leaving open Section 13.3's own tentative conclusion, and this conclusion now rests on a three-point concentration scan rather than a single value of x.

Chapter 18

A Third Profile: The Gnomonic Case

18.1 Choice of profile

Chapter 15's third named loose end asked for a third azimuthal profile beyond LAEA ($\delta(\phi) = \sec^2(\phi/2)$) and AEQD ($\delta(\phi) = \phi/\sin\phi$), to test whether Chapter 14's finding — the general boundary theorem transfers while κ_3's practical value does not — generalizes beyond two examples, and, since Chapter 16 was added later, to test whether Eq. (16.4)'s closed form is exact for LAEA alone or more broadly. This chapter uses the gnomonic projection's tangential scale factor, $\delta(\phi) = \sec\phi$, smooth and strictly increasing on $[0, \pi/2)$ with $\delta(0) = 1$, well inside this cycle's tested range ($x \leq 0.5$, against a singularity at $\pi/2 \approx 1.571$).

18.2 General boundary theorem and EVT limit

$\delta(\phi) = \sec\phi$ satisfies Theorem 10.2's hypotheses directly, giving $D = \sec x$, $\delta'(x) = D\tan x$, and $C(x, p) = \dfrac{\sin x}{2Kp\,D\tan x}$ by direct substitution into Eq. (10.3). Table 18.1 repeats Section 11.2's check for this profile: exact-quadrature ratio, a 4000-replicate Monte Carlo cross-check, and Eq. (16.4)'s prediction, at $n = 1500$.

Table 18.1: Gnomonic profile: theory mean, exact-quadrature ratio, Monte Carlo (4000 replicates) ratio, and Eq. (16.4)'s prediction, $n = 1500$. Source: `code/profile_c_gnomonic.py`, `data/profile_c_gnomonic_results.json`.

x	p	Theory mean	Exact quad. ratio	MC (4000) ratio	Equation (16.4)
0.15	1.0	0.011485	0.999319	0.999789	0.999318
0.15	2.0	0.022971	0.998638	1.007213	0.998636
0.30	1.0	0.048937	0.999272	1.017679	0.999271
0.30	2.0	0.097875	0.998544	0.979049	0.998542
0.50	1.0	0.158952	0.999148	0.988359	0.999147
0.50	2.0	0.317905	0.998298	1.003488	0.998295

The exact-quadrature ratio matches Eq. (16.4)'s prediction to five significant figures throughout, the same pattern Table 11.1 found for LAEA and AEQD; the 4000-replicate Monte Carlo column scatters around it by up to 2% per case, consistent with sampling noise of the same order documented in Section 11.2. The general boundary theorem and the EVT limit transfer to this third profile without qualification.

18.3 The second-order correction is exact here, not merely approximate

Section 16.4 left open whether Eq. (16.4)'s exact match for LAEA was a special property of that profile or held more broadly, having found only an excellent approximation (residual $\leq 2.3 \times 10^{-5}$ at $n = 1500$) for AEQD. Repeating Section 16.4's symbolic subtraction for $\delta(\phi) = \sec\phi$ gives, after simplification,

$$\text{residual} \;=\; \frac{2p}{n}\Big(\tan\tfrac{x}{2}\tan x + 1 - \sec x\Big), \tag{18.1}$$

and $\tan\frac{x}{2}\tan x + 1 - \sec x$ is identically zero for all x (direct trigonometric identity: writing $\tan\frac{x}{2} = \frac{\sin x}{1+\cos x}$, the first term equals $\frac{\sin^2 x}{(1+\cos x)\cos x}$, which combines with $1 - \sec x = \frac{\cos x - 1}{\cos x}$ to give $\frac{\sin^2 x - \sin^2 x}{(1+\cos x)\cos x} = 0$). Equation (16.4) is therefore exact for the gnomonic profile as well, with zero symbolic residual — not an approximation of the AEQD kind.

Two of the three profiles tested (LAEA, gnomonic) now give exact agreement with Eq. (16.4); only AEQD shows a nonzero, if very small, residual. This sharpens Section 16.4's open question rather than resolving it: exactness is not unique to LAEA, but it is not universal either, and Eq. (18.1)'s vanishing depends on a specific trigonometric identity satisfied by $\sec\phi$ (and, by the earlier symbolic check, by $\sec^2(\phi/2)$) that need not hold for an arbitrary profile. What distinguishes AEQD from the other two — algebraically, $\phi/\sin\phi$ is not expressible in terms of $\cos\phi$ alone, unlike $\sec\phi$ and $\sec^2(\phi/2)$, both even trigonometric functions of ϕ composed with $\cos\phi$ — is a plausible candidate explanation, but is not established as the general criterion here.

18.4 What this chapter does not test

Chapter 14's comparison of the general boundary theorem's transfer against κ_3's practical bound-rescue value required constructing κ, κ_3, and the JSD sampling pipeline for AEQD from scratch (Chapters 6 and 12's constructions, repeated). Doing the same for the gnomonic profile — refitting the two- and three-argument JSD bounds and re-running Section 13.1's stress test at $p = 20$ — is a separate, larger undertaking not carried out here; this chapter tests only the general boundary theorem, the EVT limit, and the second-order correction's generality, not κ_3's practical value for this profile.

Key Result

Chapter 15's third named loose end is partially discharged. A third azimuthal profile (gnomonic, $\delta(\phi) = \sec\phi$) confirms the general boundary theorem and EVT limit transfer without qualification, and shows Eq. (16.4) to be exact here too — not a LAEA-specific accident, since it now holds for two of three tested profiles via a shared trigonometric identity, while AEQD alone shows a small, non-zero residual. Whether κ_3's practical bound-rescue value is large or marginal for this third profile, the second half of Chapter 15's original question, is not tested here and remains open.

18.5 What the two exact cases have in common, stated precisely

The candidate explanation offered in Section 18.3 — that LAEA and gnomonic are both expressible as functions of $\cos\phi$ alone, while AEQD is not — is worth checking against the ac-

tual closed forms rather than left as an assertion. LAEA's $\delta(\phi) = \sec^2(\phi/2)$ can be rewritten, via the half-angle identity $\cos^2(\phi/2) = (1 + \cos\phi)/2$, as $\delta(\phi) = 2/(1 + \cos\phi)$: a function of $\cos\phi$ alone. Gnomonic's $\delta(\phi) = \sec\phi = 1/\cos\phi$ is trivially a function of $\cos\phi$ alone. AEQD's $\delta(\phi) = \phi/\sin\phi$, by contrast, involves ϕ itself outside of any trigonometric wrapper, not merely a function of $\cos\phi$ or $\sin\phi$; there is no algebraic substitution $\delta(\phi) = f(\cos\phi)$ for any elementary f, since $\phi/\sin\phi$ is transcendental in ϕ in a way that $2/(1+\cos\phi)$ and $1/\cos\phi$ are not. This is a genuine structural distinction, not merely a difference in how the three formulas happen to be written down: it means LAEA and gnomonic's local anisotropies are, in a precise sense, functions of position on the sphere measured through the projection's radial foreshortening alone, while AEQD's carries an explicit dependence on arc length that resists the same reduction.

This distinction is offered here as a plausible, checkable criterion, and its limits are worth stating with the same care as its content. It explains why the residual vanishes for exactly these two profiles among the three tested, but a criterion checked against two positive and one negative example is not yet a proof that every $\cos\phi$-expressible δ gives an exact match, nor that every non-$\cos\phi$-expressible δ gives a nonzero residual; both directions of that stronger claim would need an independent fourth profile to test, and this chapter does not attempt that further generalization. Chapter 22 later shows, for a different purpose, that the two classical azimuthal profiles left untested here — orthographic and stereographic — each reduce, under one of the two scale-factor conventions this cycle uses, to a closed form already tested: $\sec\phi$ (gnomonic's form) or $\sec^2(\phi/2)$ (LAEA's form). Both of those forms are functions of $\cos\phi$ alone, so this is consistent with, though it does not by itself extend, the criterion suggested here; it does not supply the independent fourth test case that would be needed to state the criterion as more than a plausible pattern across the three profiles actually computed from scratch in this cycle.

Chapter 19

The Gnomonic Case, Completed: κ_3's Practical Value for a Third Profile

19.1 The question left open

Section 18.4 stated plainly what Chapter 18 did not test: whether κ_3's practical bound-rescue value is large or marginal for the gnomonic profile, the second half of Chapter 15's third named loose end. That required refitting the two-argument (κ) and three-argument (κ_3) JSD bounds for $\delta(\phi) = \sec\phi$ from scratch — exactly Chapters 6 and 12's construction, repeated a third time — and re-running Section 13.1's $p = 20$ stress test. This chapter does that directly, using `code/gnomonic_generalization_check.py`, structurally identical to `code/aeqd_generalization_check.py` with the gnomonic central projection ($\rho(\phi) = R\tan\phi$) substituted for AEQD's equidistant one, and $\delta_{\text{gnom}}(\phi) = \sec\phi$ in place of $\delta_{\text{aeqd}}(\phi) = \phi/\sin\phi$.

Gnomonic is a useful third data point precisely because it sits, geometrically, between the other two already tested. LAEA's $\sec^2(\phi/2)$ and gnomonic's $\sec\phi$ are both even, both built from $\cos\phi$, and both diverge as $\phi \to \pi/2$; AEQD's $\phi/\sin\phi$ shares neither property, remaining finite at the pole. Section 18.3 already used this grouping to explain why Eq. (16.4)'s second-order identity is exact for LAEA and gnomonic but only approximate for AEQD. Whether that same grouping shows up in κ_3's *practical* rescue value — not merely in an exact algebraic identity, but in the size of a fitted-and-stress-tested improvement — is a separate question the algebra alone does not answer, and it is the one this chapter is built to settle.

19.2 In-range fit and the $p = 5$ check

Table 19.1 repeats Chapter 14's fit comparison (moderate range $p \le 3.0$, $n = 24$ sweep points) for the gnomonic profile, alongside the same figures already reported for LAEA (Chapter 12) and AEQD (Chapter 14).

Table 19.1: Two-argument versus three-argument (κ_3) fit span, moderate range ($p \le 3.0$), all three profiles tested in this cycle. Source: `code/gnomonic_generalization_check.py`, `data/gnomonic_generalization_results.json`; LAEA and AEQD columns repeat Chapters 12 and 14.

Profile	2-arg span	3-arg (κ_3) span	Improvement
LAEA ($\sec^2(\phi/2)$)	3.58×	1.96×	1.83×
AEQD ($\phi/\sin\phi$)	1.52×	1.49×	1.02×
Gnomonic ($\sec\phi$)	2.55×	1.56×	1.63×

The gnomonic fit is well-conditioned ($R^2 = 0.9971$, cond $= 55.3$, comparable to LAEA's 76.2 and AEQD's 89.4) and κ_3 narrows the span from 2.55× to 1.56× — a real improvement, of the same order as LAEA's, and larger than AEQD's near-negligible one.

Table 19.2 repeats Section 12.2's $p = 5$ out-of-range check.

Table 19.2: Out-of-range check at $p = 5.0$: two-argument and three-argument ratios, gnomonic profile. Source: `code/gnomonic_generalization_check.py`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	5.06×10^{-5}	5.97×	2.56×
0.30	5.35×10^{-4}	4.97×	2.06×
0.50	3.60×10^{-3}	5.01×	2.02×

At $p = 5$, κ_3 roughly halves the two-argument bound's already-modest (5–6×) underprediction to 2–2.6× — a smaller absolute failure at this point than either LAEA's (11–14× → 3.4–4.5×) or AEQD's (1.2–2.4× → 1.1–2.3×, itself already close to the two-argument value) prior results, though the three profiles are not on a common absolute scale and this ordering should not be read as a ranking of profile quality.

19.3 Stress test at $p = 20$

Table 19.3 repeats Section 13.1's stress test for the gnomonic profile.

Table 19.3: Stress test at $p = 20$: gnomonic profile, actual JSD versus two-argument and three-argument (κ_3) predictions. Source: `code/gnomonic_generalization_check.py`, `data/gnomonic_generalization_results.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	2.20×10^{-4}	543.4×	31.6×
0.30	1.45×10^{-3}	282.6×	15.8×
0.50	4.72×10^{-3}	138.4×	7.5×

The two-argument bound's failure at $p = 20$ is severe here too (138–543×), though smaller in absolute magnitude than LAEA's (1601–4249×) or, at the largest x, larger than AEQD's (6–35×). κ_3 reduces this to 7.5–31.6×, an improvement of 17–18.5× over the two-argument bound at the same point — substantial, and consistent in kind with LAEA's 44–46× improvement factor, though smaller in magnitude, and far larger than AEQD's near-flat 1.02–1.16× improvement at the same regime.

19.4 A pattern worth naming, not a proof

Section 18.3 found that Eq. (16.4) is exact, by an identical trigonometric identity, for LAEA and gnomonic, and only approximately true for AEQD — attributing this, tentatively, to LAEA and gnomonic both being even functions of ϕ composed with $\cos\phi$, unlike AEQD. This chapter's finding lines up with that grouping: κ_3's practical rescue is substantial for LAEA and gnomonic (improvement factors 1.83× and 1.63× in range, 44–46× and 17–18.5× at $p = 20$) and close to negligible for AEQD (1.02× in range, 1.02–1.16× at $p = 20$) at every comparison point checked here.

This is reported as a pattern across $n = 3$ profiles, not a mechanism. Chapter 14 already warned against over-reading a two-profile comparison; a third data point that happens to agree with the same grouping strengthens the observation but does not establish why an exact

second-order identity should predict a large κ_3 improvement, nor has any argument been constructed connecting Eq. (18.1)'s vanishing to Definition 11.2's tail statistic. No such argument is attempted here.

It is worth being explicit about what kind of caution this is, since it is easy to conflate two different worries. The first worry — that $n = 3$ is too small a sample of profiles to support a general claim about *which* profiles κ_3 will rescue — is the one raised here, and it stands regardless of how clean the LAEA/gnomonic-versus-AEQD split looks in Tables 19.1 and 19.3. The second, different worry would be that the pattern itself is not robust within these three profiles, and that is not what is being claimed: at every comparison point checked in this chapter and in Chapters 12 and 14, LAEA and gnomonic behave alike and AEQD behaves differently, with no exception. The caution is about extrapolating the grouping to a fourth, untested profile, not about whether the grouping is real for the three profiles actually examined. Chapter 22 takes up, separately, whether any fourth classical azimuthal profile could even supply that test.

Key Result

Chapter 15's third named loose end is now fully discharged. For the gnomonic profile, κ_3's practical bound-rescue value is substantial, not marginal: a $1.63\times$ narrowing of the in-range fit span, a halving of the $p = 5$ residual error, and a 17–$18.5\times$ improvement over the two-argument bound at the $p = 20$ stress point (still a genuine 7.5–$31.6\times$ absolute failure, not a rescue). Across all three profiles tested in this cycle, κ_3's practical value tracks the same LAEA/gnomonic-versus-AEQD grouping that Section 18.3's exact second-order identity check found — an observation worth stating plainly, and no more than that, given the small number of profiles tested.

Chapter 20

Is the Gnomonic Residual a Finite-n Artefact Either?

20.1 The question, asked a second time

Chapter 17 resolved, for LAEA, whether Section 13.3's noisy Monte Carlo scaling study left a genuine open question or merely an artefact of too few replicates: exact quadrature showed κ_3's own mean converges cleanly and monotonically to its EVT limit as n grows, at three concentration levels, so the $p = 20$ bound's residual failure is a structural property of the three-argument construction, not unresolved finite-n noise. Chapter 19 found a smaller but still substantial residual failure for the gnomonic profile at $p = 20$ (7.5–31.6×) without asking the same question of it. This chapter asks it, reusing `code/profile_c_gnomonic.py`'s already-verified exact-quadrature construction (Table 18.1) rather than building a new one, and, following Chapter 17's revised scope, at the same three-point concentration grid ($x = 0.15$, 0.30, 0.50) rather than at $x = 0.3$ alone.

20.2 Exact quadrature across the sample-size and concentration grid

Table 20.1: Exact-quadrature ratio of κ_3's mean to its EVT limit, gnomonic profile, $p = 20$, across $x \in \{0.15, 0.30, 0.50\}$ and $n \in \{1500, \ldots, 10^6\}$, against Eq. (16.4). Source: `code/gnomonic_kappa3_exact_convergence.py`, `data/gnomonic_kappa3_exact_convergence_multix_results.json`.

x	n	Exact ratio	$1-$ratio	$n \times (1 - \text{ratio})$	Equation (16.4)
0.15	1500	0.986551	0.013449	20.17	0.986364
0.15	5000	0.995926	0.004074	20.37	0.995909
0.15	20000	0.998978	0.001022	20.43	0.998977
0.15	100000	0.999796	0.000204	20.45	0.999796
0.15	1000000	0.999980	0.000020	20.45	0.999980
0.30	1500	0.985646	0.014354	21.53	0.985420
0.30	5000	0.995647	0.004353	21.77	0.995626
0.30	20000	0.998908	0.001092	21.84	0.998907
0.30	100000	0.999781	0.000219	21.86	0.999781
0.30	1000000	0.999978	0.000022	21.87	0.999978
0.50	1500	0.983286	0.016714	25.07	0.982947
0.50	5000	0.994915	0.005085	25.42	0.994884
0.50	20000	0.998723	0.001277	25.54	0.998721
0.50	100000	0.999744	0.000256	25.57	0.999744
0.50	1000000	0.999974	0.000026	25.58	0.999974

The pattern is the same one Table 17.1 found for LAEA: at every (x, n) pair, a clean, monotonic

rise toward 1, with Eq. (16.4)'s closed-form prediction matching to five or six significant figures throughout — essentially the same numbers as LAEA's at each matching x (e.g. $n\times(1-\text{ratio}) \to 20.9$ at $x = 0.3$ for LAEA, against 21.9 here), consistent with both profiles satisfying the same exact second-order identity (Section 18.3).

The x-dependence of the stabilization value is present here too, and more pronounced: $n \times (1 - \text{ratio})$ rises from 20.5 at $x = 0.15$ to 21.9 at $x = 0.30$ to 25.6 at $x = 0.50$, a wider spread than LAEA's $20.2 \to 20.9 \to 22.6$ over the same grid. This is consistent with Section 18.3's residual formula: both profiles carry the same x-independent leading term p/n, but the gnomonic profile's second-order coefficient, driven by $\tan x$ rather than LAEA's $\tan(x/2)$, grows faster as x increases toward the tested range's upper end, so the two profiles' stabilization values diverge more visibly at $x = 0.50$ than at $x = 0.15$.

20.3 What this settles for the gnomonic case

κ_3's own convergence to its EVT limit at $p = 20$ is not slow, stalled, or non-monotonic for the gnomonic profile at any of the three concentration levels tested: by $n = 1500$ it has already reached 98.3–98.7% of its limiting value depending on x, converging further and predictably as n grows. Chapter 19's 7.5–31.6× residual bound failure at $p = 20$, $n = 1500$ is therefore not attributable to κ_3 itself not yet having converged; it is, as for LAEA, a property of the three-argument bound's own functional form at extreme concentration, not a finite-sample artefact that more data would resolve, and this now holds uniformly across the concentration range rather than at one tested point.

Key Result

The gnomonic profile's residual κ_3-bound failure at $p = 20$ (Chapter 19) is not a finite-n convergence artefact: exact quadrature shows κ_3's own mean converging cleanly to its EVT limit at all three concentration levels tested, reaching 98.3–98.7% of it already at $n = 1500$ and 99.997–99.998% at $n = 10^6$, mirroring Chapter 17's finding for LAEA closely at each matching x ($n \times (1 - \text{ratio})$ ranging 20.5–25.6 here versus 20.2–22.6 there, with the gap between the two profiles widening as x grows). This closes the question this chapter's predecessor left implicit rather than stated, now across the full concentration range rather than a single point; nothing further is opened by it.

20.4 How the two profiles' stabilization values actually diverge

The parenthetical comparison above is worth spelling out at the same precision Table 20.1 and Chapter 17's own table report, rather than left as a qualitative "mirrors closely." Taking each profile's stabilization value at $n = 10^6$ (the largest, most converged estimate in each table) at the three matching x values: at $x = 0.15$, gnomonic gives 20.45 against LAEA's 20.23, a ratio of 1.011 (1.1% larger); at $x = 0.30$, gnomonic gives 21.87 against LAEA's 20.91, a ratio of 1.046 (4.6% larger); at $x = 0.50$, gnomonic gives 25.58 against LAEA's 22.61, a ratio of 1.131 (13.1% larger). The divergence is not merely present, as the keyresult box states, but grows roughly fourfold in relative terms from the smallest to the largest tested x ($1.1\% \to 4.6\% \to 13.1\%$), and does so monotonically rather than erratically.

This is exactly the pattern Section 18.3's residual formula predicts, not an independent finding requiring a separate explanation. Equation (18.1) shows the gnomonic and LAEA corrections share the same x-independent leading term p/n but differ in their second-order term through

$\tan x$ versus $\tan(x/2)$ respectively; since $\tan x$ grows faster than $\tan(x/2)$ as x increases over the tested range, the gap between the two profiles' second-order corrections — and hence between their stabilization values — should widen with x, and should do so smoothly rather than abruptly, exactly as the three ratios above show. The precision of this match (a monotone, accelerating divergence tracking a specific closed-form prediction, rather than a divergence of unexplained shape) is itself part of what makes Section 18.3's formula something more than a fit to the three tested points: it correctly predicts a comparative pattern — how two different profiles' residuals should relate to each other — that was not built into the formula's derivation for either profile individually.

Chapter 21

Completing the Set: Exact Convergence for AEQD

21.1 The one asymmetry left in this cycle

Chapters 17 and 20 each asked, by exact quadrature, whether κ_3's own mean converges cleanly to its EVT limit at $p = 20$ — for LAEA and gnomonic respectively, at three concentration levels ($x = 0.15, 0.30, 0.50$) — ruling out finite-n convergence as the explanation for each profile's residual bound failure. Neither chapter asked the same question of AEQD, because Chapter 14's κ_3 improvement there was already small (Table 19.1: $1.02\times$ in range, 1.02–$1.16\times$ at $p = 20$), making the question less consequential. Consequential or not, it was untested, and this chapter completes it, reusing the general, profile-independent closed form for $C(x, p)$ (Eq. (10.3), verified symbolically in `code/verify_profiles_symbolic.py` for both LAEA and AEQD) together with AEQD's $\delta(\phi) = \phi/\sin\phi$ and its closed-form derivative, handled at the removable $\phi \to 0$ singularity the same way Chapter 14's original construction already does, and at the same three-point concentration grid used for the other two profiles.

21.2 Exact quadrature across the sample-size and concentration grid

Table 21.1: Exact-quadrature ratio of κ_3's mean to its EVT limit, AEQD profile, $p = 20$, across $x \in \{0.15, 0.30, 0.50\}$ and $n \in \{1500, \ldots, 10^6\}$, against Eq. (16.4). Source: `code/aeqd_kappa3_exact_convergence.py`, `data/aeqd_kappa3_exact_convergence_multix_results.json`.

x	n	Exact ratio	$1-$ratio	$n \times (1 - \text{ratio})$	Equation (16.4)
0.15	1500	0.986726	0.013274	19.91	0.986566
0.15	5000	0.995980	0.004020	20.10	0.995970
0.15	20000	0.998992	0.001008	20.16	0.998993
0.15	100000	0.999798	0.000202	20.18	0.999799
0.15	1000000	0.999980	0.000020	20.18	0.999980
0.30	1500	0.986374	0.013626	20.44	0.986261
0.30	5000	0.995872	0.004128	20.64	0.995878
0.30	20000	0.998965	0.001035	20.71	0.998970
0.30	100000	0.999793	0.000207	20.73	0.999794
0.30	1000000	0.999979	0.000021	20.73	0.999979
0.50	1500	0.985510	0.014490	21.73	0.985512
0.50	5000	0.995605	0.004395	21.98	0.995654
0.50	20000	0.998897	0.001103	22.06	0.998913
0.50	100000	0.999779	0.000221	22.08	0.999783
0.50	1000000	0.999978	0.000022	22.08	0.999978

The pattern is the same clean, monotonic convergence found for LAEA ($n \times (1 - \text{ratio})$ ranging 20.2–22.6 across x) and gnomonic (20.5–25.6), here ranging 20.2–22.1 — the narrowest spread of the three profiles, and closest in shape to LAEA's rather than to gnomonic's wider one. The one systematic difference from those two profiles, visible at every (x, n) pair rather than only at $x = 0.3$, is that Eq. (16.4)'s prediction and the exact-quadrature value agree only to four or five significant figures here, not to the six Tables 17.1 and 20.1 showed at matching precision — consistent with Section 18.3's finding that the second-order formula is exact for LAEA and gnomonic but only an excellent approximation (residual $\leq 2.3 \times 10^{-5}$ at $n = 1500$) for AEQD; that small, already-known residual, present at all three tested x values and not concentrated at any one of them, is what separates this table from the other two, not any new discrepancy.

21.3 What this settles

κ_3's own convergence to its EVT limit at $p = 20$ is clean and monotonic for all three profiles tested in this cycle, at all three concentration levels, with no sign of stalling anywhere in the fifteen-point grid checked for any profile. For AEQD specifically, this confirms that the profile's already-small κ_3 improvement (Table 19.1) is a genuine property of the bound at this concentration, not a symptom of slow convergence that a larger sample would change — the same conclusion Chapters 17 and 20 reached for the other two profiles, now established for all three, across the same concentration range, rather than at a single shared point.

Key Result

The asymmetry this volume's revisions left open — an exact-convergence check performed for LAEA and gnomonic but not AEQD, and, until this chapter and its two predecessors, performed at a single concentration value rather than across the cycle's tested range — is closed on both counts. AEQD's κ_3 mean converges to its EVT limit exactly as cleanly as the other two profiles' at every tested x ($n \times (1 - \text{ratio})$ stabilizing in the range 20.2–22.1 here, versus 20.2–22.6 for LAEA and 20.5–25.6 for gnomonic), differing only in the already-documented, small second-order residual rather than in any convergence pathology. All three profiles tested in this cycle now carry the same exact-convergence verification across the same three-point concentration grid; nothing further is opened by it, and none is proposed.

21.4 AEQD's spread against its two companions, precisely

It is worth setting the three profiles' stabilization ranges side by side at the same x values, rather than only comparing their overall spans as the keyresult above does. At $x = 0.15$: LAEA 20.23, gnomonic 20.45, AEQD 20.18 — AEQD is, at this single point, marginally the *lowest* of the three, not merely close to LAEA. At $x = 0.30$: LAEA 20.91, gnomonic 21.87, AEQD 20.73 — again AEQD lowest, LAEA second, gnomonic highest. At $x = 0.50$: LAEA 22.61, gnomonic 25.58, AEQD 22.08 — the same ordering persists at every one of the three tested points. AEQD is therefore not merely "closest to LAEA" as Section 21.2's qualitative comparison states; it consistently sits slightly below LAEA, which itself sits below gnomonic, at all three concentration levels, a stable three-way ordering rather than a pairwise resemblance that happens to hold at one point and not others.

This ordering is worth connecting back to Section 18.3's residual analysis rather than treated as a separate, unexplained fact. Gnomonic's second-order coefficient, driven by $\tan x$, was al-

ready shown to grow fastest with x among the three profiles (Section 20.4), consistent with it having both the highest stabilization values and the widest spread across x here. AEQD's coefficient, by contrast, is close enough to LAEA's that the two profiles' curves nearly track each other, while carrying the additional small residual documented in Section 21.2 that keeps the match to four or five significant figures rather than LAEA's exact six. The three-way ordering found here is consistent with, though it does not by itself derive, that same structural picture: gnomonic's boundary rate grows fastest with x, AEQD's and LAEA's grow at similar, more moderate rates, and only AEQD carries a small additional second-order term beyond the leading behaviour the other two share exactly.

Chapter 22

The Remaining Classical Azimuthal Profiles

22.1 What is actually left untested

Three azimuthal projections have carried this cycle's local-anisotropy machinery: LAEA ($\delta(\phi) = \sec^2(\phi/2)$, Eq. (6.2)), AEQD ($\delta(\phi) = \phi/\sin\phi$, Section 14.1), and gnomonic ($\delta(\phi) = \sec\phi$, Chapter 18). Of the five classical azimuthal projections in standard use — these three plus orthographic and stereographic — only these two remain untouched. This chapter asks whether either would add a numerically new test of Theorem 10.2, rather than assuming it would and spending a chapter's worth of quadrature and Monte Carlo runs finding out.

The question is worth asking explicitly, rather than skipped, precisely because the answer is not obvious in advance. Chapters 18 and 19 treated gnomonic as a genuine third test, distinct from LAEA and AEQD, and found real differences in κ_3's practical value across the three (Section 19.4). A reader could reasonably expect the same to hold for orthographic and stereographic — two more classical projections, two more closed-form δ functions, two more chapters' worth of fitting and stress-testing. Whether that expectation is correct depends entirely on facts about these two specific projections that are checked directly below, not assumed; the alternative, that one or both turn out to be structurally degenerate or duplicative under this cycle's own framework, is exactly the kind of negative methodological result this monograph has committed, from its opening chapters onward, to reporting plainly rather than passing over.

A point easy to miss on a first pass through Chapters 6 and 18: the two chapters do not use the same δ. Equation (6.2) defines $\delta := h_\theta/h_\phi$, the ratio of tangential to radial scale factor, and this is what Chapters 6 and 10 build the closed-form CDF and κ construction from. Chapter 18's $\delta(\phi) = \sec\phi$ for gnomonic, by contrast, is h_θ alone — the tangential scale factor itself, not the ratio; gnomonic's actual ratio is $\cos\phi$, never used anywhere in this cycle. Both are legitimate inputs to Theorem 10.2, which only requires δ smooth, strictly increasing, with $\delta(0) = 1$; the theorem does not care which scale-factor quantity a profile chapter chooses to call δ. But it means "the remaining profiles" is really four candidate chapters, not two: orthographic and stereographic, each under either convention. `code/remaining_azimuthal_profiles_check.py` computes h, k, and k/h for all five classical profiles directly from their defining radial functions $f(c)$ (Snyder's $\rho = Rf(c)$ parametrization) and checks each against the closed forms already in print in this cycle; Table 22.1 reports the result.

22.2 Both remaining profiles are degenerate, for structural reasons

Stereographic is the unique azimuthal projection that is conformal: $h(c) = k(c)$ at every c, by the defining property of angle preservation, not by coincidence of this cycle's specific formulas. Under the ratio convention this cycle actually uses for its general theorem, $\delta_{\text{stereographic}}(\phi) \equiv$

Table 22.1: h, $k = f(c)/\sin c$, and k/h for all five classical azimuthal profiles at this cycle's three standard test radii. Source: `code/remaining_azimuthal_profiles_check.py`, `data/remaining_azimuthal_profiles_results.json`.

Profile	x	h	k	k/h
Gnomonic	0.50	1.298446	1.139494	0.877583
LAEA	0.50	0.968912	1.032085	1.065199
AEQD	0.50	1.000000	1.042915	1.042915
Orthographic	0.50	0.877583	1.000000	1.139494
Stereographic	0.50	1.065199	1.065199	1.000000

1 identically (confirmed in Table 22.1's last column to six decimal places at every tested x). Theorem 10.2 requires δ strictly increasing with $\delta(x) = D > 1$; a profile with $\delta \equiv 1$ has no local anisotropy for the theorem, κ, or κ_3 to act on at all. This is not a numerical near-miss to be pushed past with a larger sample: it is the exact statement that conformal projections carry no local shape distortion, which is what conformality means.

Orthographic is, dually, the unique azimuthal projection with $k(c) \equiv 1$ identically — exact preservation of scale along the parallel circle at every radius, confirmed in Table 22.1's k column. Under Chapter 18's raw-k convention, this makes orthographic degenerate in exactly the same sense stereographic is degenerate under the ratio convention: $\delta \equiv 1$, nothing to test.

Under the *other* convention, each remaining profile is well-defined but not new. Orthographic's ratio is $k/h = \sec\phi$ (Table 22.1), the identical closed form Chapter 18 already tested for gnomonic under the raw-k convention; every exact-quadrature ratio, Monte Carlo cross-check, and second-order-correction residual Chapters 17, 18 and 20 computed for gnomonic would reproduce term for term under orthographic's name, because $C(x, p)$ in Eq. (10.3) depends only on δ and its derivative at x, not on which projection produced that closed form. Symmetrically, stereographic's raw k is $\sec^2(\phi/2)$, the identical closed form Chapters 6 and 10 already tested for LAEA under the ratio convention. Running either pipeline again under the untested name would not produce a single new number; it would silently duplicate a table already in this volume under different column headings.

Table 22.1's own numbers make this duplication checkable rather than merely asserted algebraically. At $x = 0.50$, gnomonic's k/h (0.877583) matches orthographic's h to six decimal places, and gnomonic's k (1.139494) matches orthographic's k/h to the same precision — the two profiles' scale-factor pairs are literal swaps of each other, exactly as the raw-k-versus-ratio argument above predicts, not merely asymptotically or approximately equal. The same swap relationship holds between LAEA and stereographic: LAEA's k/h (1.065199) matches stereographic's h (and k, since stereographic is conformal) to six decimal places. Only AEQD's pair ($h = 1.000000$, $k/h = 1.042915$) has no swap partner among the other four in this table, consistent with AEQD being the one profile in this cycle whose defining scale factor is not built from $\cos\phi$ and does not pair off with anything else in the classical azimuthal set.

Key Result

Extending this cycle's exact-convergence and κ_3-rescue machinery to the two classical azimuthal profiles it had not yet touched — orthographic and stereographic — does not open a new test. Each is degenerate under one of the two δ-conventions this cycle has used (stereographic under the ratio convention, because it is the unique azimuthal conformal projection; orthographic under the raw-tangential-scale convention, because it uniquely preserves parallel-circle scale exactly), and under the other convention each reduces to a

closed form — $\sec\phi$ for orthographic, $\sec^2(\phi/2)$ for stereographic — already tested in this cycle under gnomonic's and LAEA's names respectively, since Eq. (10.3) depends only on that closed form and not on the projection that produced it. This is a structural fact about the five classical azimuthal profiles, verified directly rather than assumed: within this δ-based framework and this sampling model, this cycle's three tested profiles (LAEA, AEQD, gnomonic) are not merely a convenient sample of five but algebraically exhaust the classical set. A genuinely new profile-family test would require leaving azimuthal projections entirely — a cylindrical or pseudocylindrical profile, whose distortion is a function of latitude rather than of angular distance from a fixed point, and for which Section 9.4's sampling model itself would need to be rebuilt rather than merely re-parametrized. That is a materially larger undertaking than this chapter's check, is not attempted here, and is named as the one concrete direction in which "another profile" is still a genuinely open question.

Part III

Beyond the Azimuthal Family: A Latitude-Band Sampling Model

Chapter 23

Scope of Part Three

Scope

Chapter 22 named one concrete remaining direction: leaving the azimuthal family for a cylindrical or pseudocylindrical profile, whose distortion is a function of latitude rather than of angular distance from a fixed pole. This Part takes up that direction. It does not extend the azimuthal sampling model; it replaces it, because the azimuthal model's own geometry — a cap of fixed angular radius around a pole — does not describe how a cylindrical projection is used or how its distortion varies.

23.1 Why the cap model does not transfer

Every chapter since Chapter 6 has drawn sampled points from a cap of angular radius x around a fixed pole, with colatitude $\phi \in [0, x]$ distributed as $\cos\phi = 1 - V(1 - \cos x)$, $V = U^p$ (Section 9.4). This model is well matched to an azimuthal projection precisely because an azimuthal profile's local scale factors are themselves functions of colatitude from that same pole: $\delta(\phi)$ and the sampling density share one variable, angular distance from the projection's own centre.

A cylindrical or pseudocylindrical projection has no such centre. Its meridian and parallel scale factors, $h(\phi)$ and $k(\phi)$, are functions of latitude measured from the equator, constant along each parallel and independent of longitude. Points sampled from a polar cap of small angular radius are concentrated near one pole, the one place where every cylindrical projection in standard use is least representative of its own typical distortion (most are tangent or secant along low-to-middle latitudes, and several – Mercator conspicuously – degenerate at the pole itself). Reusing the cap model would not merely be inconvenient; it would test the wrong region of the sphere for this family of profiles. A new sampling model, matched to latitude rather than to angular distance from a pole, is required, and is built from first principles below rather than by analogy alone.

23.2 A latitude-band sampling model

By the north-south symmetry of every cylindrical and pseudocylindrical projection considered here, it suffices to sample the closed band of latitudes $\phi \in [0, x]$ for some $0 < x < \pi/2$, with longitude λ uniform on any interval and irrelevant to distortion; a real sampling design would draw ϕ from $[-x, x]$ symmetrically, but every $\delta(\phi)$ below is even or is evaluated only on $\phi \geq 0$, so nothing is lost restricting to the northern half.

Definition 23.1 (Latitude-band sampling model)**.** For band half-width $x \in (0, \pi/2)$ and concentration exponent $p > 0$, draw $U \sim \text{Uniform}(0, 1)$, set $V := U^p$, and set

$$\sin\phi \;=\; V \sin x. \tag{23.1}$$

Remark 23.2. On the sphere of unit radius, the area element of a band $[\phi, \phi + d\phi] \times [0, 2\pi)$ is $2\pi \cos\phi \, d\phi = 2\pi \, d(\sin\phi)$, so $\sin\phi$ uniform on $[0, \sin x]$ is exactly area-uniform sampling of the band $[0, x]$. Equation (23.1) at $p = 1$ recovers this area-uniform case, since $V = U$ is then itself uniform on $[0, 1]$ and $V \sin x$ is uniform on $[0, \sin x]$. For $p < 1$, $U^p \geq U$ typically (as $p < 1$ makes small numbers grow toward 1 faster), so V concentrates toward 1 and hence ϕ toward the boundary x; for $p > 1$, V concentrates toward 0 and ϕ toward the equator. This is the same three-regime structure — boundary-concentrated, area-uniform, centre-concentrated — that Section 9.4's cap model gave for $p < 1$, $p = 1$, $p > 1$ respectively, built here from the band's own area measure rather than transplanted from the cap.

What is genuinely new here, and what is not

Definition 23.1 is algebraically similar in shape to Section 9.4's cap model — both write a trigonometric function of the sampled angle as an affine or direct multiple of $V = U^p$ — but the trigonometric function itself changes ($\sin\phi$ against a fixed multiple of $\sin x$, rather than $\cos\phi$ against an affine map of $\cos x$), because the two models encode different sampling geometries (a band's area element is $\cos\phi \, d\phi$; a cap's is $\sin\phi \, d\phi$ referred to colatitude). The derivations that follow are carried out fully for this model rather than asserted to transfer from Chapter 10, and Section 24.4 below returns to this point once a concrete closed form is available to check.

23.3 The profile to test first

Of the classical cylindrical profiles, Mercator is conformal ($h = k$ everywhere), and Chapter 22's stereographic finding already established that a conformal profile is degenerate for this entire programme's local-anisotropy machinery: $\delta \equiv 1$, nothing for the boundary theorem to act on. The natural first non-degenerate test is the Equirectangular projection (Plate Carrée), the simplest cylindrical profile in standard use: meridians and parallels form an undistorted grid in (λ, ϕ) coordinates, with meridian scale $h(\phi) = 1$ exactly and parallel scale $k(\phi) = \sec\phi$ (the map preserves the equatorial circumference along every parallel, so a parallel at latitude ϕ, whose true circumference is $2\pi\cos\phi$ times the equator's, is stretched by $\sec\phi$ to match). The ratio

$$\delta(\phi) \;:=\; \frac{k(\phi)}{h(\phi)} \;=\; \sec\phi \tag{23.2}$$

is smooth, strictly increasing on $[0, \pi/2)$, and satisfies $\delta(0) = 1$ — exactly the hypotheses Theorem 10.2 required of an azimuthal δ, restated here for a latitude-based one. Chapter 24 derives the exact CDF of Eq. (23.2) under Definition 23.1 and the corresponding general boundary theorem, both checked numerically before anything is asserted.

What this Part does not yet do

This Part opens with one profile and one theorem, matching the discipline Chapters 10 and 11 applied to LAEA before generalizing to AEQD: derive first, generalize second, never both in the same chapter. A second cylindrical or pseudocylindrical profile, an EVT-based κ_3-style rescue statistic for this new model, and a three-argument bound fit are natural continuations of this Part and are not undertaken in this session; they are named here as open rather than silently deferred.

23.4 What transfers from the cap model's machinery, and what must be re-derived

It is useful, before Chapter 24 carries out the derivation in full, to state in advance which pieces of Chapters 10 and 11's cap-model apparatus can reasonably be expected to carry over to the band model by structural analogy, and which cannot, since conflating the two is a ready source of unearned confidence in later results.

Three things transfer at the level of proof strategy rather than formula. First, the general method of obtaining an exact CDF for δ under a monotone reparametrization of the sampled angle — express δ as a strictly increasing function of the concentration variable V, invert, and push the known CDF of $V = U^p$ through the inverse — depends only on δ being smooth and strictly increasing in the relevant angle, a hypothesis Eq. (23.2) satisfies for the band model exactly as Theorem 10.2 required of the cap model's azimuthal δ. Second, the boundary theorem's derivation, which relates the tail probability $\Pr(\delta > D - \varepsilon)$ to the local behaviour of δ's inverse near the sampling boundary, is again a statement about any such monotone δ and does not depend on whether the sampling geometry is a cap or a band. Third, the EVT bridge from a boundary theorem with a given leading rate to a specific extreme-value limit law is a general fact about regularly varying tails, established once and reused, not reproven, for each new sampling geometry this cycle introduces — as Chapter 32's discussion of the corner model's Weibull limit later makes explicit.

What does not transfer is anything downstream of the specific closed form: Eq. (10.3)'s worst-case bound constant, the numerical value of $C(x, p)$, and the resulting shape of the fitted three-argument bound are all specific to LAEA's colatitude-based δ and the cap model's $\cos\phi = 1 - V(1 - \cos x)$ reparametrization. Chapter 24 re-derives every one of these for the band model and Eq. (23.2)'s $\sec\phi$ from first principles rather than by substitution into the cap model's formulas, precisely because the disclaimerbox above in Section 23.2 already flags that the two models' trigonometric structure differs at the point where a naive relabelling would go wrong. The distinction drawn in this section — transferable proof strategy against non-transferable closed form — is the reason this Part's chapters read as full re-derivations rather than as a search-and-replace of the cap model's results, and it is the same distinction Chapter 29 draws again, more sharply, when the corner model's product structure changes not just the constants but the boundary rate's functional form itself.

Chapter 24

An Exact CDF and General Boundary Theorem for the Latitude-Band Model

24.1 The exact CDF for Equirectangular

Proposition 24.1 (Exact CDF of $\delta(\phi)$, Equirectangular, band model)**.** *Let* $D := \delta(x) = \sec x$. *Under Definition 23.1, for* $t \in [1, D]$,

$$F(t) \;=\; \Pr[\delta(\phi) \le t] \;=\; \left(\frac{\sqrt{1-t^{-2}}}{\sin x}\right)^{1/p}. \tag{24.1}$$

Proof. $\delta(\phi) \le t \iff \sec\phi \le t \iff \cos\phi \ge t^{-1} \iff \sin\phi \le \sqrt{1-t^{-2}}$ (all quantities non-negative on $\phi \in [0, \pi/2)$). Substituting $\sin\phi = V \sin x$ (Eq. (23.1)) gives $V \le \sqrt{1-t^{-2}}/\sin x$. Since $V = U^p$, $\Pr[V \le v] = v^{1/p}$ for $v \in [0,1]$, giving Eq. (24.1) directly. At $t = D = \sec x$, $1 - t^{-2} = 1 - \cos^2 x = \sin^2 x$, so the bracket equals 1 and $F(D) = 1$, as required. □

Equation (24.1) is checked against empirical CDFs from $n = 2 \times 10^5$ samples at three band half-widths and three concentration exponents, the same nine (x, p) combinations Chapter 10 used for the cap model (Table 24.1).

Table 24.1: Maximum absolute deviation between empirical and exact CDF Eq. (24.1) of $\delta(\phi) = \sec\phi$ under the band model, over 8 test points in $[1.001, D - 10^{-6}]$, $n = 2 \times 10^5$ per cell. Source: `code/band_model_general_theorem_check.py`, `data/band_model_general_theorem_results.json`.

x	p	D	Max |empirical − exact|
0.15	0.5	1.011356	0.00134
0.15	1.0	1.011356	0.00099
0.15	2.0	1.011356	0.00209
0.30	0.5	1.046752	0.00104
0.30	1.0	1.046752	0.00211
0.30	2.0	1.046752	0.00095
0.50	0.5	1.139494	0.00116
0.50	1.0	1.139494	0.00121
0.50	2.0	1.139494	0.00083

Every deviation is at or below the ≈ 0.0022 level expected from Monte Carlo noise at $n = 2\times10^5$ ($\approx 1/\sqrt{n}\cdot$const), with no systematic trend across x or p, matching the pattern Chapter 10 found for the cap model.

Key Result

Equation (24.1) is an exact, closed-form CDF for Equirectangular's local anisotropy under the latitude-band model, confirmed numerically at every tested (x, p). Like Proposition 10.1, this closed form is specific to $\delta(\phi) = \sec\phi$; Section 24.2 derives what survives for a general δ under this new model.

24.2 A general boundary theorem for the band model

Theorem 24.2 (General boundary survival rate, band model)**.** *Let $\delta : [0, x] \to [1, D]$ be smooth and strictly increasing, with $\delta(0) = 1$, $\delta(x) = D$, and $\delta'(x) > 0$. Under Definition 23.1,*

$$S(D-\varepsilon) \;=\; \Pr[\delta(\phi) > D-\varepsilon] \;\sim\; C(x,p)\,\varepsilon \qquad \text{as } \varepsilon \to 0^+, \tag{24.2}$$

with

$$C(x,p) \;=\; \frac{\cot x}{p\,\delta'(x)}. \tag{24.3}$$

Proof. Write $\delta(\phi) > D-\varepsilon \iff \phi > \phi_\varepsilon$ where $\phi_\varepsilon := \delta^{-1}(D-\varepsilon)$. A first-order expansion at x gives $D - \varepsilon \approx D - \delta'(x)(x-\phi_\varepsilon)$, so $x - \phi_\varepsilon \approx \varepsilon/\delta'(x) =: \eta$.

Under Definition 23.1, $\phi > \phi_\varepsilon \iff \sin\phi > \sin\phi_\varepsilon \iff V > V_\varepsilon := \sin\phi_\varepsilon/\sin x$, and $\Pr[V > V_\varepsilon] = 1 - V_\varepsilon^{1/p}$ as in Proposition 24.1's proof. Writing $\phi_\varepsilon = x - \eta$ and expanding $\sin\phi_\varepsilon = \sin(x-\eta) \approx \sin x - \eta\cos x$ to first order,

$$1 - V_\varepsilon = \frac{\sin x - \sin\phi_\varepsilon}{\sin x} \approx \frac{\eta\cos x}{\sin x} = \eta\cot x. \tag{24.4}$$

Since $V_\varepsilon \to 1$ as $\eta \to 0$, $1 - V_\varepsilon^{1/p} \sim \frac{1}{p}(1 - V_\varepsilon)$ by the same expansion Chapter 10 used for the cap model. Combining, $S(D-\varepsilon) \sim \eta\cot x/p \sim \varepsilon\cot x/(p\,\delta'(x))$, which is Eq. (24.2). □

Remark 24.3. Equation (24.3) has the same qualitative shape as the cap model's Eq. (10.3) — a reciprocal in p and $\delta'(x)$, times a factor depending on x alone — but the x-factor itself differs: $\cot x$ here against $1/\tan(x/2)$ there. The two are not the same function of x ($\cot x = (1-\tan^2(x/2))/(2\tan(x/2)) \neq 1/\tan(x/2)$ except in the limit $x \to 0$, where both diverge as $1/x$), so this is a genuinely different boundary rate, not Eq. (10.3) re-derived under new notation. This is expected: the two theorems share a proof strategy but are built from different area measures ($\cos\phi\,d\phi$ for the band, an affine map of $\cos\phi$ itself for the cap), and Section 23.1 gave the geometric reason those measures differ.

For Equirectangular, $\delta'(x) = \sec x\tan x$ (differentiating $\sec\phi$), giving

$$C_{\text{equirect}}(x,p) \;=\; \frac{\cot x}{p\sec x\tan x} \;=\; \frac{\cos^3 x}{p\sin^2 x}, \tag{24.5}$$

computed directly in `code/band_model_general_theorem_check.py` from Eq. (24.3) using the exact derivative, not reproduced symbolically by hand beyond the one-line simplification shown.

Table 24.2: Boundary survival rate $C(x, p)$ for Equirectangular under the band model: exact formula Eq. (24.5) versus direct least-squares estimate, $n = 2 \times 10^6$, $\varepsilon \in D \cdot \{0.002, \ldots, 0.010\}$. Source: `code/band_model_general_theorem_check.py`, `data/band_model_general_theorem_results.json`.

x	p	C_{theory}	$C_{\text{empirical}}$	Rel. error
0.15	0.5	86.575	87.665	0.013
0.15	1.0	43.288	59.354	0.371
0.15	2.0	21.644	35.443	0.638
0.30	0.5	19.968	20.200	0.012
0.30	1.0	9.984	10.591	0.061
0.30	2.0	4.992	5.397	0.081
0.50	0.5	5.881	5.969	0.015
0.50	1.0	2.941	3.014	0.025
0.50	2.0	1.470	1.518	0.032

24.3 Direct verification, and where the direct estimator is biased

Table 24.2 reports C_{theory} from Eq. (24.5) against a direct least-squares estimate at $n = 2 \times 10^6$ points, $\varepsilon \in D \cdot \{0.002, \ldots, 0.010\}$, the same design Chapter 10 used for the cap model.

The pattern here is systematic but, notably, not the same pattern found for the cap model. Chapter 10 found its largest relative errors at $p < 1$ (boundary-concentrated sampling) and small x; here the largest errors — up to 64% at $x = 0.15$ — occur at $p \geq 1$, and are worst at $p = 2.0$, the most centre-concentrated case tested, not the most boundary-concentrated one. The reason is the same finite-ε mechanism Chapter 10 identified, applied to a model where concentration now works in the opposite direction relative to the boundary layer: centre-concentrated ($p > 1$) sampling under Definition 23.1 thins out the population near $\phi = x$ precisely where the boundary theorem's linear regime lives, making a fixed finite ε grid reach that regime later, in p, than it did for the cap model.

Diagnosing, rather than dismissing, the disagreement

As in Chapter 10, this discrepancy is tested for what it is — a finite-ε artefact or a failure of Theorem 24.2 itself — by shrinking ε at fixed large n on the worst case, $x = 0.15$, $p = 2.0$ ($C_{\text{theory}} = 21.644$), rather than assumed.

Table 24.3: Convergence of the direct boundary-slope estimate to $C_{\text{theory}} = 21.644$ as $\varepsilon \to 0$, $x = 0.15$, $p = 2.0$, $n = 2 \times 10^7$. Source: `code/band_model_general_theorem_check.py`, `data/band_model_general_theorem_results.json`.

ε/D scale	Tail count	$C_{\text{empirical}}$	Rel. error
2×10^{-3}	943,138	23.314	0.077
5×10^{-4}	222,691	22.019	0.017
1×10^{-4}	43,929	21.718	0.0034
3×10^{-5}	13,007	21.435	0.0097

Relative error falls from 0.638 at the direct grid's coarsest spacing to 0.077, 0.017, and 0.0034 as ε shrinks by two further orders of magnitude, then rises slightly to 0.0097 at the smallest scale tested, where the tail count (13,007) is small enough that its own sampling noise ($\approx 1/\sqrt{13{,}007} \approx 0.9\%$) is the same order as the residual. This is the expected signature of a genuine first-order asymptotic recovering its theoretical value, with a finite-count noise floor

rather than a systematic residual at the smallest ε tested — the same qualitative resolution Section 10.3.2 found for the cap model's own worst case, on a different, newly-derived formula.

Key Result

Theorem 24.2's boundary rate Eq. (24.3) is confirmed directly for Equirectangular under the latitude-band model: the direct finite-ε estimator's worst disagreement (64% relative error, at centre-concentrated sampling, $x = 0.15$, $p = 2.0$) shrinks to 0.34% as ε narrows by two orders of magnitude, with the residual at the smallest scale tested consistent with finite tail-count noise rather than a systematic effect. The general boundary theorem transfers to the latitude-band model with a genuinely different closed form (Remark 24.3), not a relabelled copy of the cap model's result, and the specific sampling regime where the direct estimator is least reliable inverts between the two models — boundary-concentrated for the cap, centre-concentrated for the band — a difference traceable directly to which direction each model's concentration parameter moves mass relative to that model's own boundary layer.

24.4 What this opens, and what it does not yet close

This chapter derived and verified, for one cylindrical profile under one new sampling model, exactly the pair of results Chapter 10 derived and verified for one azimuthal profile under the cap model: an exact CDF, and a general boundary theorem whose rate constant is profile-specific only through $\delta'(x)$. Remark 24.3 showed directly that the two boundary-rate formulas are not the same function of x, closing the concern Section 23.1 raised in the abstract: this is new derived content, checked against simulation, not Chapter 10 copied under new symbols.

What this chapter does not do is what Chapters 11, 12 and 14 did next for the azimuthal case: define an EVT-based rescue statistic analogous to κ_3 under this model, fit a bound against it, or test a second cylindrical or pseudocylindrical profile (a natural next candidate being a pseudocylindrical profile such as sinusoidal or Mollweide, where h and k are no longer as simply related as Equirectangular's $h \equiv 1$). Each is a substantial chapter in its own right, following the same discipline — derive, then verify, then generalize — this chapter and Chapter 10 both applied, and each is named here as open rather than left unstated.

Chapter 25

Extreme-Value Convergence and κ_3 Under the Band Model

25.1 The EVT bridge transfers without a new proof

Theorem 24.2 states, for any smooth strictly increasing $\delta : [0, x] \to [1, D]$ satisfying $\delta(0) = 1$, that $S(D-\varepsilon) \sim C(x,p)\varepsilon$ as $\varepsilon \to 0^+$ under Definition 23.1. Proposition 11.1's proof, reproduced in Section 11.1, used only this one property of S — a linear survival rate at the upper endpoint — and never the cap model's own sampling mechanics; the argument is a general Gnedenko-type bridge from a linear tail rate to an exponential limit law for the rescaled maximum, not a cap-model-specific construction. It therefore transfers to the band model verbatim, with $C(x,p)$ now Eq. (24.3) in place of Eq. (10.3).

Proposition 25.1 (Exponential limit of the rescaled maximum, band model)**.** *Under the hypotheses of Theorem 24.2, as $n \to \infty$,*

$$n\big(D - M_n\big) \xrightarrow{d} \text{Exponential}\big(C(x,p)\big), \tag{25.1}$$

with $C(x,p)$ as in Eq. (24.3) and $M_n := \max_{i \le n} \delta(\phi_i)$.

Proof. Identical to Proposition 11.1's proof sketch, substituting Eq. (24.2) for Eq. (10.2): $\Pr[n(D - M_n) > t] = (1 - S(D - t/n))^n \to e^{-Ct}$ for fixed $t > 0$ as $n \to \infty$, using only $S(D - t/n) \sim Ct/n$. □

Definition 25.2 (The κ_3 statistic, band model)**.** For a sample of n points drawn under Definition 23.1, with sample maximum local anisotropy $\kappa_2^{\max} := \max_{i \le n} \delta(\phi_i)$ and worst-case bound $D = \delta(x)$,

$$\kappa_3 \;:=\; n\big(D - \kappa_2^{\max}\big), \tag{25.2}$$

exactly Definition 11.2, unmodified: the statistic itself does not depend on which sampling model or which theorem supplies its limiting distribution, only Proposition 25.1's rate constant $C(x,p)$ does.

By Proposition 25.1, $\kappa_3 \xrightarrow{d} \text{Exponential}(C(x,p))$, so $\mathbb{E}[\kappa_3] \to 1/C(x,p)$, a finite non-degenerate limit under the band model exactly as under the cap model.

25.2 Simulation check: Equirectangular

Table 25.1 reports, for Equirectangular ($\delta(\phi) = \sec\phi$, the only band-model profile with an established general boundary theorem so far, Chapter 24), at three band half-widths and $p \in \{1, 2\}$, $n = 1500$: the exact-quadrature mean of $n(D - M_n)$ in ϕ-space (using Definition 23.1's closed-form ϕ-CDF, $G(\phi) = (\sin\phi / \sin x)^{1/p}$, in place of Section 11.2's cap-model

Table 25.1: Exact-quadrature and Monte Carlo (4000 replicates) means of $n(D - M_n)$ against theory, Equirectangular, band model, $n = 1500$. Source: `code/band_evt_kappa3.py`, `data/band_evt_kappa3_results.json`.

x	p	Theory mean	Exact quad.	Ratio (quad)	MC (4000)	Ratio (MC)
0.15	1.0	0.023101	0.023069	0.9986	0.023169	1.0029
0.15	2.0	0.046203	0.046075	0.9972	0.045254	0.9795
0.30	1.0	0.100163	0.100010	0.9985	0.099959	0.9979
0.30	2.0	0.200325	0.199716	0.9970	0.201503	1.0059
0.50	1.0	0.340078	0.339423	0.9981	0.339497	0.9985
0.50	2.0	0.680156	0.677543	0.9962	0.702809	1.0333

G), the theoretical mean $1/C(x, p)$, their ratio, and an independent 4000-replicate Monte Carlo mean.

The exact-quadrature ratio sits at 0.9962–0.9986 across all six cases, a small finite-n shortfall of the same 0.1–0.4% order Section 11.2 found for LAEA and AEQD under the cap model, though slightly larger here and growing mildly with p rather than staying essentially flat; the independent 4000-replicate Monte Carlo column agrees with it within its own sampling noise at every case tested.

Table 25.2 reports the accompanying KS statistics against the full Exponential($C(x, p)$) CDF, same design.

Table 25.2: KS statistic against Exponential($C(x, p)$), Equirectangular, band model, $n = 1500$, 4000 replicates. Source: `code/band_evt_kappa3.py`, `data/band_evt_kappa3_results.json`.

x	p	KS statistic	p-value
0.15	1.0	0.0142	0.393
0.15	2.0	0.0176	0.165
0.30	1.0	0.0067	0.994
0.30	2.0	0.0138	0.423
0.50	1.0	0.0226	0.033
0.50	2.0	0.0275	0.005

Four of six cases are non-significant at conventional levels; the two rejections both occur at $x = 0.50$, the largest band half-width tested, where Section 24.3 already found the general boundary theorem's own finite-ε estimator converging more slowly than at smaller x (Table 24.2's pattern, worst at centre-concentrated sampling). This is consistent with the same underlying approach to the asymptotic regime being marginally slower at larger x under this model, not with a distinct failure of Proposition 25.1 itself, and is reported rather than smoothed over.

Key Result

Proposition 25.1 extends the EVT bridge from the cap model to the band model by direct reuse of Proposition 11.1's proof, requiring no new argument beyond substituting Eq. (24.2)'s rate. For Equirectangular, it is confirmed at $n = 1500$ by exact quadrature (ratio 0.996–0.999), an independent Monte Carlo check, and a KS test (non-significant in four of six cases, both rejections at the largest band width tested). Definition 25.2 defines κ_3 for the band model with no modification to Definition 11.2's formula; Chapter 26 tests whether it rescues the JSD bound the way Chapter 12 showed it does under the cap model.

25.3 The shortfall depends on p as well as on x

Table 25.1's six exact-quadrature ratios are worth reading along both axes of the design, not only along x as the main discussion above emphasizes. Fixing x and comparing $p = 1$ against $p = 2$: at $x = 0.15$, the ratio moves from 0.9986 to 0.9972; at $x = 0.30$, from 0.9985 to 0.9970; at $x = 0.50$, from 0.9981 to 0.9962. The finite-n shortfall from 1 therefore roughly doubles when p doubles from 1 to 2, consistently at all three band half-widths, in addition to growing with x at fixed p. Both trends point the same direction — greater concentration of sampling mass (whether through a larger p or a wider band) leaves the exact-quadrature estimate marginally further from its asymptotic limit at the fixed sample size $n = 1500$ used throughout Table 25.1 — which is the same qualitative direction Section 24.3's finite-ε estimator convergence showed for the underlying boundary theorem itself, and is consistent with both results reflecting the same finite-sample approach to the same asymptotic regime rather than two unrelated effects.

The KS test's two rejections, both at $x = 0.50$ (Table 25.2), sharpen this same pattern rather than introducing a new one: the p-values at $x = 0.50$ fall to 0.033 ($p = 1$) and 0.005 ($p = 2$), both well below the four non-rejected cases' range of 0.165 to 0.994, and the ordering — $p = 2$'s p-value roughly a sixth of $p = 1$'s at the same x — matches the ordering already seen in the exact-quadrature ratios at that x. This internal consistency across two different diagnostics (a mean-ratio comparison and a full-distribution KS test), both showing the same x-then-p ordering of severity, is part of why Section 25.2's original discussion attributes the two rejections to a slower approach to the asymptotic regime at the widest, most concentrated sampling settings, rather than treating them as two isolated and independently-explained anomalies.

Chapter 26

Does κ_3 Rescue the Bound Under the Band Model?

26.1 A JSD pipeline for the band model

Chapters 6 and 12's JSD pipeline — sample points, project them with a candidate coordinate map, compute the Jensen–Shannon divergence between the histograms of pairwise geodesic and pairwise planar distances, and fit a power-law bound in log space — has no cap-model-specific step beyond the sampling call and the coordinate map itself, so it is reused here with both replaced: sampling from Definition 23.1 in place of the cap model, and Equirectangular coordinates $(X, Y) = (\lambda, \phi)$ (Chapter 24) in place of LAEA's. Longitude $\lambda \sim \text{Uniform}(-\pi, \pi)$ is drawn independently of ϕ at each replicate, matching Section 23.2's remark that longitude is irrelevant to distortion and may be drawn from any interval; geodesic distance uses the standard latitude-based great-circle formula, $\cos\gamma = \sin\phi_1 \sin\phi_2 + \cos\phi_1 \cos\phi_2 \cos(\lambda_1 - \lambda_2)$, in place of Chapter 6's colatitude-based one. κ is the sampling-weighted mean of $\delta(\phi) = \sec\phi$ over the sampled points, and κ_3 is Definition 25.2, both computed identically to their cap-model counterparts.

26.2 The three-argument fit

Fitting, on the identical 24-point $(x, p \le 3.0)$ grid Chapters 6 and 12 used,

$$\log \text{JSD} = a \log(\kappa - 1) + b \log(D - 1) + d \log \kappa_3 + c \tag{26.1}$$

gives $a = -0.105$, $b = 0.328$, $d = -0.162$, $c = -1.683$, $R^2 = 0.8707$, condition number 43.7, against a two-argument fit on the same grid of $a = 0.153$, $b = -0.098$, $c = -1.505$, $R^2 = 0.8356$.

Table 26.1: Two-argument versus three-argument (κ_3) fit, band model, Equirectangular, identical 24-point grid, $p \le 3.0$. Source: `code/band_kappa3_bound.py`, `data/band_kappa3_bound_results.json`.

Fit	R^2	In-range span	Condition number
Two-argument (κ, D)	0.8356	1.18×	—
Three-argument (κ, D, κ_3)	0.8707	1.18×	43.7

Two features distinguish this fit from Chapter 12's cap-model result before any interpretation is attempted. First, the absolute R^2 values are markedly lower here (0.84, 0.87 against 0.97, 0.99): log JSD spans only -1.92 to -1.51 across the whole 24-point grid (Table 26.2 below), a much narrower target than the cap model's multi-order-of-magnitude range, so a modest absolute residual translates into a larger fraction of the (small) explained variance. Second, the in-range span is unchanged by adding κ_3 ($1.18\times$ either way): the two-argument fit was

already tight in this range, unlike Chapter 6's cap-model fit, which had a $3.58\times$ span for κ_3 to narrow.

Table 26.2: Representative rows of the JSD sweep, band model, Equirectangular ($N = 1500$ points, 8 replicates per cell). Source: `data/band_kappa3_bound_results.json`.

x	p	κ	JSD	κ_3
0.15	0.2	1.0081	0.1652	0.0045
0.15	3.0	1.0016	0.1466	0.0666
0.50	0.2	1.0951	0.2204	0.0655
0.50	3.0	1.0182	0.1493	0.9796

Why the range is narrow: a scale mismatch this chapter did not remove

Section 23.2 draws λ from a full $[-\pi, \pi]$ interval regardless of the band half-width $x \in \{0.15, 0.30, 0.50\}$, so the sampled region is a thin, full-circumference ring rather than a small patch shrinking uniformly in every direction the way Section 9.4's cap does as $x \to 0$. The JSD is consequently dominated by the fixed, large longitudinal extent rather than by x's own shrinkage, which compresses the JSD's dependence on (x, p) into a narrow band regardless of the profile or fit used. This is a genuine structural difference between the two sampling geometries, not an artefact of the fitting procedure, and it is the reason Chapter 24's honesty section did not promise this chapter's fit would reproduce Chapter 12's dramatic in-range narrowing.

26.3 Out-of-range tests: $p = 5$ and $p = 20$

Table 26.3: Out-of-range test at $p = 5$: actual JSD versus two-argument and three-argument (κ_3) predictions, band model, Equirectangular. Source: `data/band_kappa3_bound_results.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	0.14585	$1.17\times$	$1.10\times$
0.30	0.14652	$1.09\times$	$1.03\times$
0.50	0.14712	$1.04\times$	$0.98\times$

Table 26.4: Stress test at $p = 20$ (outside the fitted range and outside Table 26.3's test point), band model, Equirectangular. Source: `data/band_kappa3_p20_stress.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	0.14435	$1.42\times$	$1.18\times$
0.30	0.14451	$1.32\times$	$1.09\times$
0.50	0.14471	$1.25\times$	$1.04\times$

At both out-of-range test points, adding κ_3 moves every ratio closer to 1, and the direction of the residual error is the same at $p = 5$ and $p = 20$ (mild overprediction by the two-argument bound at small x, easing toward slight underprediction by the three-argument bound at large x), not a sign flip of the kind Section 7.2 found for $\kappa_2^{\max}$. This is a genuine, if modest, rescue, in contrast to Chapter 12's cap-model finding along two distinct axes: there is no severe underprediction crisis here for κ_3 to solve (the two-argument bound's worst ratio, $1.42\times$ at $p = 20$, is nowhere

near the cap model's 36–97× residual failure at $p = 20$, Section 13.1), and the improvement, while consistent in direction at every one of the six test cells across both stress points, is a narrowing of a mild error rather than a rescue from a severe one.

26.4 What the narrow-range result implies for choosing a rescue statistic in practice

Sections 26.1 to 26.3 together settle the technical question this chapter opened with, but a practitioner deciding whether to invest in computing κ_3 for a specific projection and sampling geometry needs a criterion, not just this chapter's and Chapter 12's two data points. The comparison between them supplies one directly.

The two chapters differ in exactly one relevant respect for this purpose: the underlying sampling geometry's own range of achievable JSD values at fixed concentration parameter p. The cap model's angular radius x shrinks the sampled region uniformly in every direction as $x \to 0$, so its JSD range spans several orders of magnitude across the tested grid; the band model's longitude sweep stays fixed at the full $[-\pi, \pi]$ interval regardless of x (Table 26.2's disclaimerbox), so its own JSD range is compressed into roughly half an order of magnitude by comparison. κ_3's absolute correction, in both cases, is bounded above by the size of the residual the two-argument bound already leaves unexplained, and that residual is itself bounded above by the target's own dynamic range: there is more room for a rescue statistic to matter when the quantity being bounded ranges widely than when it does not. This is exactly the mechanism Section 28.1 states in general terms once both this chapter's and Chapter 27's results are in hand.

The practical rule this suggests, stated here as a heuristic rather than a proven theorem, is: before computing κ_3 (or any analogous higher-order correction) for a new sampling geometry, first check the two-argument bound's own residual range across the geometry's natural concentration and extent parameters. A geometry whose JSD already varies by orders of magnitude at fixed extent, as the cap model's does, is a geometry where a tail-sensitive correction is likely to earn its computational cost; a geometry whose extent is fixed and whose concentration parameter alone drives a narrow JSD range, as the band model's is, is one where the same correction, while never harmful in the direction it moves the estimate (Section 26.3), is unlikely to be worth pursuing before the two-argument bound is judged adequate on its own terms. Chapter 32's account of the corner model's own bound-rescue value, left for a future pipeline run rather than computed in this session, is exactly the case this heuristic would flag as worth checking directly rather than assuming from either of this Part's two precedents.

Key Result

Under the band model, adding κ_3 to the JSD bound raises R^2 from 0.8356 to 0.8707 without narrowing the already-tight in-range span ($1.18\times$ either way), and moves every out-of-range test ratio closer to 1 at both $p = 5$ (six cells, 1.04–$1.17\times \to 0.98$–$1.10\times$) and the more extreme $p = 20$ stress point (1.25–$1.42\times \to 1.04$–$1.18\times$), consistently in the correcting direction. The rescue is real but modest, because the band model's own JSD range is narrow to begin with (Table 26.2's disclaimerbox): this is a genuinely different outcome from Chapter 12's dramatic cap-model narrowing, not a weaker replication of it, and the difference is traced directly to the two sampling geometries' different relationship between their concentration parameter and their fixed spatial extent.

Chapter 27

A Second Cylindrical Profile: Where $h \neq 1$

27.1 Choice of profile, and a substitution disclosed rather than made silently

Section 24.4 named a second cylindrical or pseudocylindrical profile as open, suggesting sinusoidal or Mollweide as candidates precisely because, unlike Equirectangular, their meridian scale factor h is not identically 1. Both, however, have a property that Equirectangular does not: away from the central meridian, their meridians are curved, so their meridian scale factor depends on longitude as well as latitude. Concretely, for any pseudocylindrical map of the form $X = f(\phi)\lambda$, $Y = g(\phi)$, the meridian direction at fixed $\lambda_0 \neq 0$ has plane displacement $(f'(\phi)\lambda_0,\, g'(\phi))\, d\phi$, giving $h(\phi, \lambda_0) = \sqrt{f'(\phi)^2\lambda_0^2 + g'(\phi)^2}$, which reduces to a function of ϕ alone only at $\lambda_0 = 0$. Definition 23.1 and every construction built on it through Chapter 26 assume δ is a function of ϕ alone; testing sinusoidal or Mollweide honestly, at longitudes away from the central meridian, would require rebuilding that machinery around a two-variable $\delta(\phi, \lambda)$, a materially larger undertaking than this chapter attempts and not carried out here.

This chapter tests the Lambert cylindrical equal-area projection instead: $X = \lambda$, $Y = \sin\phi$. Its meridians are straight ($f(\phi) \equiv 1$ in the notation above, so there is no shear at any longitude, and δ is exactly a function of ϕ alone, matching Definition 23.1's hypotheses without modification), while its meridian scale factor is $h(\phi) = \cos\phi \not\equiv 1$ — the specific property named as open, obtained without the shear complication. This substitution is disclosed here as a substitution, not offered as sinusoidal or Mollweide under another name.

Proposition 27.1 (Local anisotropy, Lambert cylindrical equal-area)**.** *For $X = \lambda$, $Y = \sin\phi$, the parallel scale factor is $k(\phi) = \sec\phi$ and the meridian scale factor is $h(\phi) = \cos\phi$, giving*

$$\delta(\phi) \;:=\; \frac{k(\phi)}{h(\phi)} \;=\; \sec^2\phi. \tag{27.1}$$

Proof. Along a parallel ($d\phi = 0$): $dX = d\lambda$, sphere distance $= \cos\phi\, d\lambda$, so $k = 1/\cos\phi = \sec\phi$. Along a meridian ($d\lambda = 0$): $dY = \cos\phi\, d\phi$, sphere distance $= d\phi$, so $h = \cos\phi$. The ratio is Eq. (27.1). □

$\delta(\phi) = \sec^2\phi$ is smooth and strictly increasing on $[0, \pi/2)$ with $\delta(0) = 1$, satisfying Theorem 24.2's hypotheses; $\delta'(\phi) = 2\sec^2\phi\tan\phi$.

27.2 Exact CDF and general boundary theorem

Proposition 27.2 (Exact CDF of $\delta(\phi) = \sec^2\phi$, band model)**.** *Let* $D := \sec^2 x$. *For* $t \in [1, D]$,

$$F(t) \;=\; \left(\frac{\sqrt{1-t^{-1}}}{\sin x}\right)^{1/p}. \tag{27.2}$$

Proof. $\delta(\phi) \le t \iff \sec^2\phi \le t \iff \cos^2\phi \ge t^{-1} \iff \sin\phi \le \sqrt{1-t^{-1}}$. Substituting $\sin\phi = V\sin x$ and $\Pr[V \le v] = v^{1/p}$ gives Eq. (27.2) directly, exactly as Proposition 24.1's proof did for $\delta(\phi) = \sec\phi$, with t^{-2} replaced by t^{-1} throughout because $\delta = \sec^2\phi$ rather than $\sec\phi$. At $t = D$, $1 - t^{-1} = 1 - \cos^2 x = \sin^2 x$, so the bracket is 1 and $F(D) = 1$. □

Table 27.1 checks Eq. (27.2) against empirical CDFs at the same nine (x, p) cells Chapter 24 used for Equirectangular.

Table 27.1: Maximum absolute deviation between empirical and exact CDF Eq. (27.2), Lambert cylindrical equal-area, band model, $n = 2 \times 10^5$. Source: `code/band_second_profile_lambertcea.py`, `data/band_second_profile_lambertcea_results.json`.

x	p	D	Max |empirical − exact|
0.15	0.5	1.02282	0.00154
0.15	1.0	1.02282	0.00074
0.15	2.0	1.02282	0.00183
0.30	0.5	1.09554	0.00096
0.30	1.0	1.09554	0.00191
0.30	2.0	1.09554	0.00110
0.50	0.5	1.29847	0.00106
0.50	1.0	1.29847	0.00140
0.50	2.0	1.29847	0.00121

Every deviation is at the same $\lesssim 0.002$ Monte Carlo noise level Table 24.1 found for Equirectangular. Applying Theorem 24.2 (reused without modification, since it was proved for arbitrary δ) with $\delta'(x) = 2\sec^2 x \tan x$ gives

$$C_{\text{lcea}}(x, p) \;=\; \frac{\cot x}{2p\sec^2 x \tan x} \;=\; \frac{\cos^4 x}{2p\sin^2 x}, \tag{27.3}$$

a different function of x from Eq. (24.5)'s $\cos^3 x/(p\sin^2 x)$ — an extra power of $\cos x$ and a factor of 2, both traceable directly to $\delta'(x)$'s extra factor of $2\sec^2 x \tan x$ against $\sec x \tan x$ — so this is again new derived content, not Eq. (24.5) relabelled.

The worst case ($x = 0.15$, $p = 2.0$, 18.9%) is diagnosed the same way Section 24.3 diagnosed its own worst case, by shrinking ε at fixed large n rather than assumed to be a theorem failure.

Relative error falls from 0.189 to 0.038, 0.006, and 0.004–0.002 (the last two consistent with each other within finite tail-count noise) as ε narrows by two further orders of magnitude, the same asymptotic-recovery signature Section 24.3 found for Equirectangular's own worst case.

Key Result

Proposition 27.2 and Eq. (27.3) extend Chapter 24's exact-CDF-and-boundary-theorem pair to a second cylindrical profile with $h \neq 1$, chosen (Lambert cylindrical equal-area, in

Table 27.2: Boundary rate $C(x, p)$, Lambert cylindrical equal-area: exact formula Eq. (27.3) versus direct estimate, $n = 2 \times 10^6$. Source: `code/band_second_profile_lambertcea.py`.

x	p	C_{theory}	$C_{\text{empirical}}$	Rel. error
0.15	0.5	42.802	43.213	0.010
0.15	1.0	21.401	24.029	0.123
0.15	2.0	10.700	12.721	0.189
0.30	0.5	9.538	9.612	0.008
0.30	1.0	4.769	4.929	0.034
0.30	2.0	2.385	2.471	0.036
0.50	0.5	2.581	2.616	0.014
0.50	1.0	1.290	1.309	0.015
0.50	2.0	0.645	0.654	0.013

Table 27.3: Convergence of the direct boundary-slope estimate to $C_{\text{theory}} = 10.700$ as $\varepsilon \to 0$, $x = 0.15$, $p = 2.0$, $n = 2 \times 10^7$. Source: `data/band_second_profile_eps_convergence.json`.

ε/D scale	Tail count	$C_{\text{empirical}}$	Rel. error
2×10^{-3}	454,257	11.103	0.038
5×10^{-4}	110,129	10.767	0.006
1×10^{-4}	21,801	10.657	0.004
3×10^{-5}	6,577	10.717	0.002

place of the shear-complicated sinusoidal or Mollweide originally suggested, Section 27.1) specifically to isolate that one property. Both results are confirmed numerically at the same standard of evidence Chapter 24 applied to Equirectangular, including an ε-convergence diagnostic of the same worst-case cell.

27.3 EVT limit and κ_3 for this profile

Proposition 25.1 applies to $\delta(\phi) = \sec^2 \phi$ without modification, giving $n(D - M_n) \xrightarrow{d}$ Exponential$(C_{\text{lcea}}(x, p))$. Table 27.4 repeats Section 25.2's check for this profile.

Table 27.4: Exact-quadrature and Monte Carlo (4000 replicates) means of $n(D - M_n)$, Lambert cylindrical equal-area, band model, $n = 1500$. Source: `data/band_second_profile_lambertcea_results.json`.

x	p	Theory mean	Exact quad.	Ratio (quad)	MC (4000)	Ratio (MC)
0.15	1.0	0.046727	0.046662	0.9986	0.046651	0.9984
0.15	2.0	0.093454	0.093195	0.9972	0.091797	0.9823
0.30	1.0	0.209691	0.209358	0.9984	0.212425	1.0130
0.30	2.0	0.419381	0.418054	0.9968	0.414705	0.9889
0.50	1.0	0.775033	0.773388	0.9979	0.781689	1.0086
0.50	2.0	1.550067	1.543502	0.9958	1.507019	0.9722

The exact-quadrature ratio (0.9958–0.9986) matches Table 25.1's Equirectangular pattern closely, at every corresponding (x, p) cell, within a few parts in 10^4; KS statistics against the full exponential law range 0.0120–0.0217 (non-significant in five of six cases, the one rejection at $x = 0.30, p = 1.0, p = 0.046$), consistent with the same finite-replicate testing noise Section 25.2

reported.

27.4 Practical κ_3 bound-rescue test

Rebuilding Chapter 26's JSD pipeline with $\kappa = \text{mean}(\sec^2\phi)$, $D = \sec^2 x$, and Lambert cylindrical equal-area coordinates $(X, Y) = (\lambda, \sin\phi)$ in place of Equirectangular's gives a two-argument fit ($a = 0.153$, $b = -0.098$, $c = -1.505$, $R^2 = 0.8356$) and three-argument fit ($a = -0.105$, $b = 0.328$, $d = -0.162$, $c = -1.683$, $R^2 = 0.8707$, condition number 43.7) essentially matching Table 26.1's Equirectangular coefficients to two or three significant figures, and an out-of-range $p = 5$ test (Table 27.5) matching Table 26.3 nearly exactly.

Table 27.5: Out-of-range test at $p = 5$, Lambert cylindrical equal-area, band model. Source: `data/band_second_profile_lambertcea_results.json`.

x	Actual JSD	Two-arg ratio	Three-arg (κ_3) ratio
0.15	0.14585	1.17×	1.11×
0.30	0.14652	1.09×	1.03×
0.50	0.14714	1.04×	0.98×

This near-exact match across two profiles that differ in δ, in D, and in the plane coordinate map is itself informative, not a coincidence to note in passing: Section 26.1's disclaimerbox already identified the mechanism, the longitude range $\lambda \in [-\pi, \pi]$ fixed independently of x, as what compresses the JSD's dependence on (x, p) into a narrow band for Equirectangular; the same mechanism operates identically here, since it depends only on the sampling geometry (band half-width and full longitude sweep) and not on which profile's δ or plane map is used. Where the azimuthal-model comparison (Chapter 14) found κ_3's practical value to be profile-dependent because the two-argument bound's baseline severity differed sharply between LAEA and AEQD, the band-model comparison across Equirectangular and Lambert cylindrical equal-area finds the opposite: both profiles share the same longitude-dominated baseline severity, so both share essentially the same modest rescue.

Key Result

The general boundary theorem, the EVT limit, and κ_3's modest practical bound-rescue value all transfer to a second cylindrical profile with $h \neq 1$ (Lambert cylindrical equal-area) at the same standard of evidence applied to Equirectangular, discharging Section 24.4's third named open item, with the disclosed substitution of Section 27.1. Unlike the azimuthal case's LAEA-versus-AEQD comparison, this second profile does not reveal profile-dependence in κ_3's practical value; the reason is structural, traced to Section 26.1's longitude-dominated sampling geometry, which affects any band-model profile equally and is a property of the sampling model this chapter did not vary, not of the two profiles compared.

Chapter 28

Part Three: Closing Accounting

Section 24.4 named three items explicitly open at the end of Chapter 24: an EVT-based κ_3-style rescue statistic for the band model, a three-argument bound fit against it, and a second cylindrical or pseudocylindrical profile. All three are now closed.

- Chapter 25 showed the cap model's EVT bridge (Proposition 11.1) transfers to the band model by direct reuse of its proof, defined κ_3 for the band model without modifying Definition 11.2's formula, and confirmed the resulting exponential limit law for Equirectangular by exact quadrature, Monte Carlo, and KS test.
- Chapter 26 built a JSD pipeline for the band model and showed κ_3 raises the three-argument fit's R^2 from 0.836 to 0.871 and moves every out-of-range test ratio toward 1 at both $p = 5$ and a $p = 20$ stress point, a real but modest rescue traced to the band model's own narrow JSD range rather than to any deficiency in κ_3 itself.
- Chapter 27 tested a second cylindrical profile with $h \neq 1$ (Lambert cylindrical equal-area, substituted for the shear-complicated sinusoidal or Mollweide originally suggested, with the substitution disclosed rather than made silently) and found the exact CDF, general boundary theorem, EVT limit, and κ_3's practical bound-rescue value all transfer at the same standard of evidence, with the added finding that κ_3's practical value here is not profile-dependent the way it was for the azimuthal LAEA/AEQD comparison — a structural consequence of the band model's fixed longitude sweep, not of the two cylindrical profiles compared.

28.1 The bound-rescue result in full

The headline number in the second item above — R^2 rising from 0.836 to 0.871 — understates how the rescue actually behaves across the sweep, and it is worth stating the mechanism rather than only the summary statistic, since the same pattern recurs for the corner model in Chapters 32 and 36 and is easy to mistake for a coincidence if seen only once.

At the $p = 20$ stress point, Chapter 26's pipeline recorded the two- and three-argument fit ratios at three representative x values on the band model's own JSD scale: at $x = 0.15$ the two-argument bound overshoots the true value by a factor of 1.417, which κ_3 brings down to 1.182; at $x = 0.3$ the corresponding pair is $1.318 \to 1.094$; at $x = 0.5$ it is $1.255 \to 1.040$. The pattern is monotone and one-directional across all three points — κ_3 never makes the fit worse anywhere it was tested — but the absolute improvement shrinks as x grows, because the band model's own JSD range at fixed p is narrow, so there is less headroom for any rescue statistic, κ_3 included, to close. This is the concrete content behind the qualifier "real but modest" used in the summary above: the rescue is genuine at every tested point, and its size is bounded above by a property of the band model's sampling geometry, not by a weakness specific to κ_3's functional form.

A second point worth making explicit is what did *not* have to be re-derived. Definition 11.2's formula, fixed for the cap model in Chapter 12, was reused for the band model without modification; only the sampling distribution feeding it changed. That the same functional form transfers across two structurally different sampling geometries (radial concentration around a pole for the cap model, a one-dimensional latitude margin for the band model) without needing a new free parameter is itself part of what Chapter 25 closed, and it is the reason Chapter 26 could proceed directly to a bound-fit test rather than first needing to justify a new definition.

28.2 Why Lambert cylindrical equal-area, and what that choice costs

The substitution named in the third item deserves one further sentence of motivation beyond the disclosure already given in Section 27.1. Sinusoidal and Mollweide were the profiles originally named because they are the cylindrical family's standard shear-bearing examples; both were set aside for this Part specifically because their scale factor depends on longitude as well as latitude, which is exactly the two-variable structure Chapter 29 goes on to build machinery for in Part Four. Lambert cylindrical equal-area was chosen instead precisely because it keeps $h \neq 1$ (so the test is not vacuous) while keeping δ a function of latitude alone (so Chapter 23's one-variable band model still applies without modification). The cost of that choice is exactly the item recorded in the disclaimer below: nothing in this Part speaks to whether κ_3's profile-independence finding survives once genuine two-variable shear is present, and Part Four's corner model is the place that question is finally taken up directly, rather than answered by extrapolation from this Part's necessarily narrower test.

What remains genuinely open

Section 27.1's substitution is itself a scope boundary, not a resolved question: whether κ_3's practical value depends on profile choice for a genuinely shear-bearing pseudocylindrical projection (sinusoidal, Mollweide) at longitudes away from the central meridian — which would require extending Definition 23.1's machinery to a two-variable $\delta(\phi, \lambda)$ — is not tested anywhere in this Part and is named here, once, as the concrete direction left for further work, rather than folded into the closing list above as though it had been addressed by proxy.

28.3 This Part in relation to Part Two

It is worth setting this Part's overall outcome explicitly alongside Chapter 15's, since the two Parts address structurally parallel questions — does an EVT-based rescue statistic improve a JSD bound fit, for a given sampling geometry and profile — and reach different answers for reasons that are themselves informative rather than merely coincidental.

Part Two's cap-model result was dramatic and uneven: a clean, well-conditioned rescue in the moderate-concentration range (Chapter 12's span narrowing from $3.58\times$ to $1.96\times$), a real but partial correction at $p = 5$, and an honest, quantified failure at $p = 20$ (Chapter 13). This Part's band-model result, by contrast, is modest and uniform: a real but small rescue at every tested point, with no regime in which either the two-argument bound fails badly or κ_3's correction is dramatic. Section 28.1 traced this to the band model's narrower JSD range rather than to any weakness in κ_3 itself, and Section 28.2 traced the choice of Lambert cylindrical equal-area, rather than a more strongly shear-bearing alternative, to keeping this Part's one-variable machinery applicable at all.

The two Parts together therefore make a point neither could make alone: κ_3's value as a correction is not a fixed property of the statistic, but depends on the sampling geometry's own dynamic range in a way this cycle can now state as a general pattern rather than as an observation specific to one profile. A single Part, reporting only the cap model's dramatic rescue or only the band model's modest one, would risk a reader generalizing incorrectly in either direction — assuming κ_3 always helps a great deal, or assuming it rarely helps much. Reporting both, with the mechanism identified in each case, is what allows Section 32.1's later account of the corner model's quadratic boundary rate to be read as a third data point in an emerging pattern, rather than as an isolated new phenomenon requiring its own explanation from scratch.

Part IV

The Two-Variable Shear Profile

Chapter 29

Part Four: The Two-Variable Shear Profile

Scope

Chapter 28's closing disclaimer named exactly one concrete item as genuinely open: whether κ_3's practical rescue value, and the boundary/EVT machinery beneath it, extend to a genuinely shear-bearing pseudocylindrical profile — one whose local anisotropy δ depends on *two* variables, latitude ϕ and longitude λ, rather than on ϕ alone. This Part takes up that item directly, for the profile named there: the Sinusoidal (Sanson–Flamsteed) projection away from its central meridian. It does not attempt a fully general two-variable boundary theory for an arbitrary $\delta(\phi, \lambda)$; Section 29.3 states precisely what is and is not covered.

29.1 The profile: Sinusoidal away from the central meridian

The Sinusoidal projection is defined by $x_{\text{map}} = \lambda \cos\phi$, $y_{\text{map}} = \phi$ (radius-one sphere, λ and ϕ in radians). Its parallel scale is $k \equiv 1$ (parallels map to their true length), and its meridian scale is

$$h(\phi, \lambda) = \sqrt{1 + \lambda^2 \sin^2\phi}, \tag{29.1}$$

obtained directly from the map equations' partial derivatives rather than quoted from a secondary source. Both were re-derived and checked symbolically in `code/verify_sinusoidal_scale_factors.py`, which confirms $k \equiv 1$, confirms Eq. (29.1), and confirms the identity $h^2 - 1 = \lambda^2 \sin^2\phi$ used throughout this Part. On the central meridian ($\lambda = 0$) or the equator ($\phi = 0$), $h = 1 = k$: the projection is locally conformal there, and $\delta := h/k = h$ reduces to the degenerate value 1 studied nowhere further. Away from both, $\delta(\phi, \lambda) = \sqrt{1 + \lambda^2 \sin^2\phi} > 1$ genuinely depends on both arguments — the shear this Part's title refers to, and the reason Chapters 24 to 27's one-variable machinery does not apply to it as it stands.

29.2 Why a new sampling model is needed, and what it assumes

Definition 23.1's latitude-band model draws ϕ from a band and treats λ as irrelevant to distortion, which was correct for every profile Part Three tested ($h \equiv 1$ or $k \equiv 1$ with the other a function of ϕ alone). It is not correct here: Eq. (29.1) makes δ genuinely a function of the pair (ϕ, λ), so a sampling model must specify how both are drawn jointly.

Definition 29.1 (Corner-shear sampling model)**.** For latitude half-width $x \in (0, \pi/2)$, longitude half-width $\lambda_{\max} > 0$, and concentration exponents $p, q > 0$, draw $U_1, U_2 \sim \text{Uniform}(0, 1)$

independently, set $V := U_1^p$, $W := U_2^q$, and set

$$\sin\phi \;=\; V \sin x, \qquad |\lambda| \;=\; \lambda_{\max}\, W, \tag{29.2}$$

with the sign of λ drawn independently and uniformly from $\{-1, +1\}$ (irrelevant to δ, since Eq. (29.1) depends on λ only through λ^2).

A disclosed modeling choice, not a derived fact

Definition 29.1 samples ϕ and λ *independently*. This is a genuine simplification, stated here rather than left implicit: a real gazetteer of place names is not laid out as an independent product of a latitude margin and a longitude margin, and Section 29.3 returns to this. The independence assumption is what makes an exact closed-form analysis possible in this chapter; it is not claimed to be the only reasonable sampling design, only a tractable and clearly-stated one, in the same spirit as Definition 23.1 and Section 9.4's cap model each having been one explicit choice among others.

Remark 29.2. At $p = 1$, V is uniform on $[0, 1]$ and $\sin\phi$ is uniform on $[0, \sin x]$, area-uniform sampling of the latitude band exactly as in Definition 23.1. At $q = 1$, W is uniform on $[0, 1]$ and λ is uniform on $[-\lambda_{\max}, \lambda_{\max}]$. General p, q give the same boundary-concentrated / uniform / centre-concentrated three-regime structure in each variable separately, now controlled independently in latitude and longitude.

29.3 The reduction that makes this tractable

Substituting Eq. (29.2) into Eq. (29.1),

$$\delta(\phi, \lambda) \;=\; \sqrt{1 + \lambda^2 \sin^2\phi} \;=\; \sqrt{1 + a^2 Z^2}, \qquad a := \lambda_{\max} \sin x, \quad Z := VW. \tag{29.3}$$

δ depends on the pair (V, W) only through their product $Z = VW \in [0, 1]$: a genuine structural fact about this specific multiplicative shear form, not a modeling simplification. It is what reduces a nominally two-dimensional boundary problem to a one-dimensional one in Z, and Chapter 30 exploits it directly. The joint maximum $D := \sqrt{1 + a^2}$ is attained only at the single point $\phi = x$, $|\lambda| = \lambda_{\max}$ (i.e. $V = W = 1$, $Z = 1$) — an isolated corner of the sampling rectangle, not an edge or curve as in every profile Part Two and Part Three studied, and Chapter 31 shows this changes the boundary rate's functional form, not merely its constant.

What this reduction does and does not generalize

Equation (29.3)'s clean reduction to a single product variable $Z = VW$ is specific to Eq. (29.1)'s form, $h^2 = 1 + \lambda^2 \sin^2\phi$, in which λ and $\sin\phi$ enter only as a product inside the square root. The Mollweide projection's shear, mentioned alongside Sinusoidal in Chapter 28's open item, has a different closed form not of this product type, and nothing in this Part is claimed to transfer to it automatically; Mollweide is named again, and left open again, in Chapter 32.

What this Part does not do

This Part derives and verifies an exact CDF for δ under Definition 29.1 (Chapter 30), a general boundary theorem and its EVT bridge (Chapter 31), matching in kind, though not in every technical detail, what Chapters 10 and 11 did for the cap model and Chapters 24 and 25 did for the band model. It does not build a κ_3-analogue JSD bound-fit against real

gazetteer data for this model, the way Chapters 12 and 26 did for their respective models: that would require a full pipeline run against real coordinate data under a two-margin sampling design and is not undertaken in this session. It is named here as open, not silently folded into what follows.

29.4 Why the isolated corner is the structurally new case, not merely a harder instance of the old one

It is worth being explicit, before Chapters 30 and 31 carry out the technical derivation, about why this Part is titled around a genuinely new profile type rather than presented as a routine extension of Part Three's one-variable machinery to a second variable. The distinction matters because it sets the reader's expectations for what kind of result to look for.

Every boundary theorem before this Part — Chapter 10's for the cap model, Chapter 24's for the band model — concerns a supremum of δ attained along a one-dimensional edge of the sampling region: the cap's boundary circle at angular radius x, or the band's boundary parallel at latitude x. In both cases, the tail event $\{\delta > D - \varepsilon\}$ for small ε is, to leading order, an event about a single scalar margin (colatitude, or latitude) being within $O(\varepsilon)$ of its extreme value, and the corresponding probability is linear in ε because that margin's sampling density is bounded and positive at the boundary. Equation (29.3)'s reduction to the product variable $Z = VW$ changes this qualitatively rather than quantitatively: the supremum of δ is attained at a single *point* (Section 29.3's corner $V = W = 1$), not along a curve, and the tail event $\{Z > 1 - \eta\}$ near that point requires *both* margins simultaneously near their own extremes, which is a two-dimensional constraint collapsing to a vanishing area rather than a one-dimensional constraint collapsing to a vanishing length.

This is precisely the geometric fact Section 32.1 returns to once Chapter 31's theorem is in hand and can be stated with its quadratic rate attached; the point made here, in advance of that derivation, is only that the corner model's qualitative novelty — point supremum against curve supremum — was visible already from Eq. (29.3)'s algebra, before any boundary probability was computed, and is the reason this Part could not have been written as a mechanical two-variable generalization of Chapter 24's proof template even in principle.

Chapter 30

An Exact CDF for the Two-Variable Shear Profile

30.1 The distribution of a product of two power-law variables

Every closed-form CDF in this cycle prior to this chapter came from a single power-law variable: the azimuthal cap model and the latitude-band model of Parts I and II both reduce, ultimately, to a one-dimensional sampling exponent p acting on a single radial or latitudinal coordinate. Sinusoidal's two-variable corner-shear model, introduced in Chapter 29 to capture the shear that appears once both ϕ and λ vary away from the central meridian and equator simultaneously, has no such reduction available: $\delta(\phi, \lambda)$ depends on the product of two independently sampled coordinates, each with its own concentration exponent (p for ϕ, q for λ). An exact distributional result for this model therefore requires a genuinely new derivation, not a relabelling of a one-variable result already in print, which is what this section supplies.

By Section 29.3, $\delta(\phi, \lambda) \le t \iff Z \le z(t) := \sqrt{t^2 - 1}/a$ for $t \in [1, D]$, so an exact CDF for δ reduces to an exact CDF for $Z = VW$, $V = U_1^p$, $W = U_2^q$, U_1, U_2 independent uniforms.

Proposition 30.1 (Exact CDF of a product of two independent power-law variables)**.** *Let V, W be independent with $\Pr[V \le v] = v^\alpha$, $\Pr[W \le w] = w^\beta$ on $[0, 1]$ $(\alpha, \beta > 0)$, and $Z := VW$. For $z \in [0, 1]$,*

$$F_Z(z) \;=\; \begin{cases} \dfrac{\alpha z^\beta - \beta z^\alpha}{\alpha - \beta}, & \alpha \neq \beta, \\[2ex] z^\alpha\big(1 - \alpha \ln z\big), & \alpha = \beta. \end{cases} \tag{30.1}$$

Proof. Condition on $V = v$: for $v \le z$, $\Pr[W \le z/v \mid V = v] = 1$ since $z/v \ge 1$; for $v > z$, $\Pr[W \le z/v \mid V = v] = (z/v)^\beta$. With density $f_V(v) = \alpha v^{\alpha-1}$,

$$F_Z(z) = \int_0^z \alpha v^{\alpha-1}\, dv + \int_z^1 \left(\frac{z}{v}\right)^\beta \alpha v^{\alpha-1}\, dv = z^\alpha + \alpha z^\beta \int_z^1 v^{\alpha-1-\beta}\, dv.$$

For $\alpha \neq \beta$, the remaining integral is $(1 - z^{\alpha-\beta})/(\alpha - \beta)$, giving $F_Z(z) = z^\alpha + \frac{\alpha}{\alpha-\beta}(z^\beta - z^\alpha)$, which simplifies to Eq. (30.1)'s first case by collecting the z^α terms. The $\alpha = \beta$ case follows either by the same integral with $\int_z^1 v^{-1} dv = -\ln z$, or as the $\beta \to \alpha$ limit of the first case by L'Hopital's rule in β; both routes give $z^\alpha(1 - \alpha \ln z)$, matching the classical density $-\alpha^2 z^{\alpha-1} \ln z$ for a product of two i.i.d. Beta$(\alpha, 1)$ variables at $\alpha = 1$ (the familiar $F(z) = z - z \ln z$ for a product of two independent standard uniforms). □

Key Result

Equation (30.1) gives, via $t = \sqrt{1 + a^2 z^2}$, an exact closed-form CDF for $\delta(\phi, \lambda)$ under Definition 29.1, for any concentration exponents p, q and any band/window half-widths

$x, \lambda_{\max}$: with $\alpha = 1/p$, $\beta = 1/q$, $a = \lambda_{\max} \sin x$,

$$F_\delta(t) \;=\; F_Z\left(\frac{\sqrt{t^2-1}}{a}\right), \qquad t \in [1, D],\; D = \sqrt{1+a^2}. \tag{30.2}$$

30.2 Numerical verification

Equation (30.1) is checked against empirical CDFs of $Z = VW$ from $n = 2 \times 10^6$ Monte Carlo draws, at nine (p, q) combinations spanning the equal-exponent case ($\alpha = \beta$, the branch requiring the logarithmic limit) and unequal exponents in both orders, evaluated at nine test points $z \in [0.05, 0.95]$ per cell.

Table 30.1: Maximum and mean absolute deviation between empirical and exact CDF Eq. (30.1) of $Z = VW$, over 9 test points $z \in [0.05, 0.95]$, $n = 2 \times 10^6$ per cell. Expected Monte Carlo noise scale $1/\sqrt{n} \approx 0.00071$. Source: `code/corner_shear_exact_cdf.py`, `data/corner_shear_exact_cdf_results.json`.

p	q	$\alpha = 1/p$	$\beta = 1/q$	Max |dev|	Mean |dev|
0.5	0.5	2.0	2.0	0.00029	0.00015
1.0	1.0	1.0	1.0	0.00021	0.00008
2.0	2.0	0.5	0.5	0.00016	0.00007
0.5	1.0	2.0	1.0	0.00041	0.00023
0.5	2.0	2.0	0.5	0.00022	0.00009
1.0	2.0	1.0	0.5	0.00026	0.00011
1.0	0.5	1.0	2.0	0.00030	0.00010
2.0	0.5	0.5	2.0	0.00058	0.00015
2.0	1.0	0.5	1.0	0.00033	0.00013

Every maximum deviation, including the three equal-exponent cells that exercise the logarithmic branch of Eq. (30.1) rather than the ratio branch, is at or below the ≈ 0.00071 level expected from Monte Carlo noise at $n = 2 \times 10^6$, with no systematic trend across (p, q) and no gap in accuracy between the two branches of Eq. (30.1).

Key Result

Proposition 30.1 is confirmed numerically at nine (p, q) combinations, both branches, to within Monte Carlo noise. Composed with Eq. (30.2), this gives an exact, closed-form CDF for the Sinusoidal projection's local anisotropy under the two-variable corner-shear model — the first exact distributional result of this Part, and the one Chapter 31's asymptotic theorem is built on and checked against.

30.3 Reading the two branches together

Equation (30.1)'s two cases are not two independent formulas that happen to agree at $\alpha = \beta$ by construction; the proof derives the equal-exponent case as the $\beta \to \alpha$ limit of the general one, so the two branches are a single analytic function and the equal-exponent row of Table 30.1 is exercising the same underlying expression as its six unequal-exponent neighbours, not a separately-verified special case standing apart from them. This matters for how

Table 30.1's numbers should be read: the equal-exponent cells ($p = q \in \{0.5, 1, 2\}$) show maximum deviations of 0.00029, 0.00021, and 0.00016 respectively, each comfortably inside the six unequal-exponent cells' range of 0.00022 to 0.00058, with no discontinuity in accuracy at the boundary where the logarithmic case takes over from the ratio case. Had the two branches been independently coded, a systematic accuracy gap at exactly $\alpha = \beta$ would have been the natural symptom of an implementation error at the limit; its absence here is itself part of the verification, alongside the direct comparison against Monte Carlo.

The asymmetric cells also show a mild but consistent pattern worth naming rather than leaving implicit: the two orderings of the same (p, q) pair are not numerically identical in Table 30.1, since V and W play different roles (one is folded through p, the other through q, and Eq. (30.2)'s outer map is not symmetric in the two exponents once $x \neq \lambda_{\max}$ generally), and the largest single deviation in the table, 0.00058 at $(p, q) = (2.0, 0.5)$, is roughly double the smallest, 0.00016 at $(p, q) = (2.0, 2.0)$ — still an order of magnitude below anything that would indicate a formula error, and consistent with ordinary Monte Carlo variability across nine independently drawn cells rather than a systematic asymmetry in Eq. (30.1) itself. This closed-form CDF is exactly what Chapter 31 differentiates near its upper endpoint to get the boundary rate, and what Chapter 37 later re-tests against genuine GeoNames coordinates rather than the synthetic two-margin sample used throughout this chapter.

It is worth flagging, in advance of that later use, exactly what this chapter has and has not established. Proposition 30.1 is an exact statement about the model random variable $Z = VW$ under Definition 29.1's specific sampling assumptions — independent power-law margins in ϕ and λ — verified here against synthetic draws from that same model. It says nothing yet about whether real settlement coordinates are well approximated by independent power-law margins in the first place; that is a separate empirical question, addressed only later, once Chapter 31 has built the asymptotic boundary theory this exact CDF feeds into. Keeping the two questions — is the closed form correct given the model, and does the model fit real data — separate is deliberate: Table 30.1's agreement with Monte Carlo speaks only to the first.

Chapter 31

A Corner-Type Boundary Theorem and Its EVT Bridge

31.1 Why the boundary rate changes shape here

Every boundary theorem so far in this monograph — Theorem 10.2 for the cap model, Theorem 24.2 for the band model — found $S(D-\varepsilon) \sim C\varepsilon$, linear in ε, because in each case the maximum of δ is attained along an entire boundary curve (the rim of the cap, or the top of the latitude band) reached by a single sampled coordinate. Under Definition 29.1, D is attained only at the single point $(\phi, \lambda) = (x, \lambda_{\max})$, i.e. $V = W = 1$ simultaneously: an isolated corner of the sampling rectangle, not an edge. Reaching within ε of D now requires *both* V and W to be simultaneously close to 1, and Section 29.3 showed this is governed entirely by $Z = VW$ being close to 1 — a strictly more restrictive event than either margin alone being close to 1, which is exactly why the rate below turns out to be quadratic rather than linear.

31.2 The corner boundary theorem

Theorem 31.1 (Corner-type boundary survival rate)**.** *Under Definition 29.1, with* $\alpha = 1/p$, $\beta = 1/q$, $a = \lambda_{\max} \sin x$, $D = \sqrt{1+a^2}$,

$$S(D-\varepsilon) \;=\; \Pr[\delta(\phi,\lambda) > D-\varepsilon] \;\sim\; K\,\varepsilon^2 \qquad \text{as } \varepsilon \to 0^+, \qquad K \;=\; \frac{\alpha\beta D^2}{2a^4}. \tag{31.1}$$

Proof. Write $\delta > D-\varepsilon \iff Z > 1-\eta$ for the corresponding η (via Eq. (29.3)). Since $D^2 - t^2 = (D-t)(D+t) \approx 2D\varepsilon$ for $t = D-\varepsilon$ near D, and $D^2 - t^2 = a^2(1-Z^2) \approx 2a^2\eta$ for $Z = 1-\eta$ near 1, matching the two gives $\eta \approx (D/a^2)\varepsilon$.

For the survival of Z itself, Proposition 30.1 gives $F_Z'(z) = \alpha\beta(z^{\beta-1} - z^{\alpha-1})/(\alpha-\beta)$ for $\alpha \neq \beta$ (the $\alpha = \beta$ case is the limit of what follows). Writing $z = 1-\eta$ and expanding $z^k \approx 1 - k\eta + \frac{1}{2}k(k-1)\eta^2$ to second order,

$$z^{\beta-1} - z^{\alpha-1} \;\approx\; (\alpha-\beta)\eta \;-\; (\alpha-\beta)\frac{\alpha+\beta-3}{2}\eta^2,$$

so $F_Z'(1-\eta) \approx \alpha\beta\,\eta + O(\eta^2)$: the density of Z vanishes linearly as $z \to 1^-$ (in contrast to a generic bounded density, because $F_Z'(1) = 0$ identically, a direct consequence of both marginal densities $\alpha v^{\alpha-1}$ and $\beta w^{\beta-1}$ needing to act together for Z to approach 1). Integrating, $S_Z(1-\eta) = \int_0^\eta F_Z'(1-s)\,ds \approx \frac{1}{2}\alpha\beta\,\eta^2$.

Combining with $\eta \approx (D/a^2)\varepsilon$: $S(D-\varepsilon) = S_Z(1-\eta) \sim \frac{1}{2}\alpha\beta(D/a^2)^2\varepsilon^2 = K\varepsilon^2$, which is Eq. (31.1). □

Remark 31.2. Equation (31.1)'s quadratic exponent is not a special feature of the Sinusoidal profile's particular constants; it traces directly to the corner geometry of Definition 29.1 combined

with δ depending on (V, W) only through their product. Any profile whose local anisotropy reduces, under an independent-margin sampling model, to a smooth strictly increasing function of a product $Z = VW$ of two independent power-law variables would show the same ε^2 rate with the same proof; Eq. (30.1)'s vanishing derivative at $z = 1$ is the general mechanism, not an accident of this one profile.

31.3 Numerical verification: the boundary rate

Table 31.1 reports, for three $(x, \lambda_{\max}, p, q)$ combinations, the fitted log–log slope of $S(D - \varepsilon)$ against ε over five geometrically-spaced ε values at $n = 2 \times 10^7$ direct samples — a slope near 2 confirms Eq. (31.1)'s exponent directly, independent of any theoretical K.

Table 31.1: Direct finite-ε verification of the corner boundary rate, $n = 2 \times 10^7$, $\varepsilon \in D \cdot \{0.00125, \ldots, 0.02\}$. Source: `code/corner_shear_boundary_evt.py`, `data/corner_shear_boundary_results.json`.

x	$\lambda_{\max}$	p	q	K_{theory}	Fitted log–log slope
0.30	0.8	1.0	1.0	168.998	2.241
0.30	0.8	0.5	2.0	168.998	2.225
0.50	1.2	2.0	1.0	3.037	2.065

All three fitted slopes are close to the theoretical exponent 2 (against 1, which every earlier boundary theorem in this monograph found), with the modest excess above 2 attributable to the same finite-ε curvature effect Chapters 10 and 24 both diagnosed for their own (linear) rates. Table 31.2 isolates this for the first case, shrinking ε at fixed large n.

Table 31.2: Convergence of the direct boundary-slope estimate of K to $K_{\text{theory}} = 168.998$ as $\varepsilon \to 0$, $x = 0.30$, $\lambda_{\max} = 0.8$, $p = q = 1.0$, $n = 2 \times 10^7$. Source: as Table 31.1.

ε/D	Tail count	$K_{\text{empirical}}$	Rel. error
0.0200	3,040,179	359.906	1.130
0.0100	475,294	225.067	0.332
0.0050	101,506	192.266	0.138
0.0025	23,757	180.996	0.065
0.00125	5,754	174.381	0.032

Relative error falls monotonically from 1.13 to 0.032 as ε shrinks across five geometrically-spaced scales, the same qualitative first-order-asymptotic signature Section 24.3 found for the band model's own worst case, and consistent with Eq. (31.1) being the genuine leading-order behavior rather than an artefact of the coarsest grid point.

Key Result

Theorem 31.1's corner-type boundary rate, quadratic in ε rather than linear, is confirmed both by a direct log–log slope fit (2.07–2.24 against a theoretical exponent of 2) and by convergence of the fitted rate constant K toward K_{theory} as $\varepsilon \to 0$ (relative error falling from 1.13 to 0.032 over five scales). This is a genuinely different functional form from every boundary rate found earlier in this monograph, traced in Section 31.1 to the isolated-corner geometry of Definition 29.1, not merely a different constant re-derived under new notation.

31.4 The EVT bridge: a Weibull, not exponential, limit

Proposition 31.3 (Rayleigh-type limit of the rescaled maximum)**.** *Under the hypotheses of Theorem 31.1, as* $n \to \infty$*,*

$$\sqrt{n}\,(D - M_n) \xrightarrow{d} T, \qquad \Pr[T > t] = \exp(-Kt^2), \qquad t \geq 0, \tag{31.2}$$

a Weibull distribution of shape 2 *(a half-normal/Rayleigh-type law), with* K *exactly Eq. (31.1)'s rate constant.*

Proof sketch. $\Pr[\sqrt{n}(D-M_n) > t] = \Pr[M_n < D - t/\sqrt{n}] = F(D-t/\sqrt{n})^n = (1 - S(D-t/\sqrt{n}))^n$. By Eq. (31.1), $S(D - t/\sqrt{n}) \sim Kt^2/n$ for fixed t as $n \to \infty$, so $(1 - Kt^2/n + o(1/n))^n \to \exp(-Kt^2)$: the same Gnedenko-type argument Proposition 11.1 used, with $\sqrt{n}$ rescaling in place of n because the boundary rate here is order ε^2 rather than ε^1. □

A genuinely different rescaling, not the same κ_3 reused

Definition 11.2's $\kappa_3 := n(D - \kappa_2^{\max})$ is built for a linear boundary rate and is not the correct statistic here: under Eq. (31.1)'s quadratic rate, $n(D - M_n) \to 0$ in probability (over-rescaled) rather than converging to a non-degenerate law. The corner analogue is $\kappa_3^{\text{corner}} := \sqrt{n}\,(D - M_n)$, converging instead to the Weibull(2) law of Eq. (31.2). This is stated explicitly here because Section 11.3's original definition, unmodified, would silently give a degenerate statistic if applied to this model without noting the rescaling must change.

31.5 Numerical verification: the EVT bridge

Table 31.3 checks Eq. (31.2) by exact quadrature: composing Eq. (30.2) with $(\cdot)^n$ gives $\Pr[M_n < D - t/\sqrt{n}]$ exactly, with no Monte Carlo noise, at three sample sizes chosen so that $\sqrt{n}\,(D-1)$ — the point beyond which the finite range $[1, D]$ itself, rather than the asymptotic law, would dominate — comfortably exceeds the range where $\exp(-Kt^2)$ has already decayed to 10^{-8}.

Table 31.3: Exact-quadrature check of Eq. (31.2): maximum absolute deviation of $\Pr[\sqrt{n}(D - M_n) > t]$ from $\exp(-Kt^2)$ over $t \in [0, t_{\max}]$, and ratio of the quadrature mean of $\sqrt{n}(D - M_n)$ to the theoretical Rayleigh mean $\sqrt{\pi/(4K)}$. Source: `code/corner_shear_evt_quadrature_v2.py`, `data/corner_shear_evt_quadrature_v2_results.json`.

$x, \lambda_{\max}, p, q$	K	n	Max |dev|	Mean (num.)	Mean (theory)	Ratio
$0.30, 0.8, 1, 1$	168.998	14,343	0.00626	0.06749	0.06817	0.9900
		57,375	0.00311	0.06778	0.06817	0.9943
		229,501	0.00155	0.06793	0.06817	0.9964
$0.30, 0.8, 0.5, 2$	168.998	14,343	0.00545	0.06756	0.06817	0.9911
		57,375	0.00271	0.06782	0.06817	0.9948
		229,501	0.00135	0.06795	0.06817	0.9967
$0.50, 1.2, 2, 1$	3.037	25,677	0.00646	0.50393	0.50850	0.9910
		102,710	0.00322	0.50616	0.50850	0.9954
		410,843	0.00160	0.50728	0.50850	0.9976

At every one of the three cases, the maximum survival-function deviation shrinks by very close to a factor of 2 each time n quadruples (consistent with an $O(1/\sqrt{n})$ finite-n correction to the

limit law), and the mean ratio rises monotonically toward 1 ($0.990 \to 0.996$), the same small, consistent finite-n shortfall pattern Section 11.2 found for the linear-rate cap and band models.

Table 31.4 adds an independent stochastic check: direct Monte Carlo simulation of $\kappa_3^{\text{corner}} = \sqrt{n}(D - M_n)$ over 4000 replicates at $n = 2000$, entirely independent of the exact-quadrature machinery above.

Table 31.4: Direct Monte Carlo check of $\mathbb{E}[\kappa_3^{\text{corner}}]$ against the theoretical Rayleigh mean $\sqrt{\pi/(4K)}$, $n = 2000$ samples per replicate, 4000 replicates. Source: `code/corner_shear_boundary_evt.py`, `data/corner_shear_evt_mc_results.json`.

$x, \lambda_{\max}, p, q$	K	MC mean	Theory mean	Ratio
$0.30, 0.8, 1, 1$	168.998	0.06636	0.06817	0.9735
$0.30, 0.8, 0.5, 2$	168.998	0.06659	0.06817	0.9769
$0.50, 1.2, 2, 1$	3.037	0.49213	0.50850	0.9678

The independent Monte Carlo ratios (0.968–0.977) agree with the exact-quadrature ratios' direction and rough magnitude (both are small, consistent finite-n/finite-replicate shortfalls from the asymptotic mean, not a discrepancy between the two verification routes), matching the pattern Section 11.2 established for the linear-rate models.

Key Result

The EVT bridge for the corner-shear model, Proposition 31.3, is confirmed by two independent routes: noise-free exact quadrature at three sample sizes (survival-function deviation shrinking as n grows, mean ratio 0.990–0.996), and a genuinely stochastic 4000-replicate Monte Carlo check (mean ratio 0.968–0.977). The correct rescaling is $\sqrt{n}$, not n, and the correct limit law is Weibull of shape 2, not exponential — both direct, verified consequences of Theorem 31.1's quadratic rather than linear boundary rate, closing Chapter 28's named open item for the boundary-theorem and EVT layer of this programme's machinery.

Chapter 32

Part Four: Closing Accounting

Chapter 28's disclaimer named one concrete open item: whether κ_3's practical value, and the boundary/EVT machinery beneath it, depend on profile choice for a genuinely shear-bearing pseudocylindrical projection at longitudes away from the central meridian, which would require extending the machinery to a two-variable $\delta(\phi, \lambda)$. This is now closed, at the specific scope stated below and in Section 29.3.

- Section 29.1 re-derived, symbolically, the Sinusoidal projection's scale factors ($k \equiv 1$, $h = \sqrt{1 + \lambda^2 \sin^2 \phi}$), confirming the two-variable dependence named as untested in Chapter 28.
- Definition 29.1 built an explicit, disclosed two-margin sampling model (independent latitude and longitude concentration exponents p, q), and Section 29.3 showed the resulting δ depends on the two margins only through their product $Z = VW$ — a genuine structural fact about this profile's shear form, not an assumption.
- Chapter 30 derived an exact closed-form CDF for δ under this model (Proposition 30.1, Eq. (30.2)), confirmed numerically at nine (p, q) combinations to within Monte Carlo noise.
- Chapter 31 derived and verified a boundary theorem of a genuinely new type: quadratic, not linear, in ε (Theorem 31.1), traced to the isolated-corner geometry this model introduces; derived and verified the corresponding EVT bridge, a Weibull(2) rather than exponential limit law under $\sqrt{n}$ rather than n rescaling (Proposition 31.3); and named the correctly-rescaled corner analogue of κ_3 explicitly, rather than reusing Definition 11.2's formula unmodified where it would silently degenerate.

32.1 Why the corner geometry forces a quadratic rate

The quadratic boundary rate named in the fourth item is the single most structurally distinctive result in this Part, and it is worth being explicit about why it appears here and nowhere earlier in the monograph, rather than leaving it as an assertion to be taken on faith. Every boundary theorem before this Part — the cap model's (Chapter 10) and the band model's (Chapter 24) alike — concerns a tail event defined by a single scalar margin crossing a threshold, and in both cases the tail probability is, to leading order, linear in ε because the margin's density is bounded and positive at the threshold. The corner model's tail event is different in kind: because δ depends on the *product* $Z = VW$ of two independent margins (Section 29.3), the tail $\{Z > z\}$ for z near its supremum is an event about *both* margins simultaneously being extreme, and the measure of a product-tail region shrinks quadratically, not linearly, in the distance from the corner — geometrically, the same reason the area of a small square corner cut from a rectangle scales as the square, not the first power, of the cut's size. Theorem 31.1 makes this exact rather than heuristic, and Chapter 31's verification data confirm it directly: at the tested case $(x, \lambda_{\max}, p, q) = (0.3, 0.8, 1, 1)$, the relative error between the theoretical quadratic

coefficient and its empirical estimate falls from 1.13 at the coarsest tested ε to 0.032 at the finest, a monotone convergence consistent with the claimed rate and not with a linear one, which the same data would have shown as a residual error plateauing rather than continuing to fall as $\varepsilon \to 0$.

The Weibull(2) limit law under $\sqrt{n}$ rescaling named in the same item is the distributional consequence of this quadratic rate, by the same extreme-value argument Chapter 25 used to derive the band model's exponential limit — only here the regularly-varying tail index is 2 rather than 1, which is exactly what changes an exponential limit into a Weibull one and an n-rescaling into a $\sqrt{n}$-rescaling. Chapter 31's Monte Carlo replication at three sample sizes for the case $(x, \lambda_{\max}, p, q) = (0.3, 0.8, 1, 1)$ found the ratio of empirical to theoretical mean climbing from 0.993 at $n = 4599$ toward 1 as n grows to 73597, the finite-n convergence pattern expected of a correctly identified limit law rather than a misspecified one.

32.2 This Part's place in the cycle's sequence of sampling models

Three sampling models now span this monograph: the cap model (Section 9.4), matched to the azimuthal family's shared pole; the band model (Definition 23.1), matched to the cylindrical family's shared dependence on latitude alone; and the corner model (Definition 29.1), matched to a shear-bearing profile whose local anisotropy genuinely depends on two coordinates. Each was introduced only once its predecessor was shown, on structural grounds rather than by assumption, not to transfer to the next profile family under study — Section 23.1's argument for why the cap model does not describe a cylindrical projection's typical distortion, and Section 29.4's argument for why the corner model's point supremum is a different kind of object from either predecessor's curve supremum.

Reading the three models as a single progression rather than as three independent constructions clarifies what has actually been established by this Part's closing. It is not that a single, maximally general sampling and boundary theory has been built and instantiated three times; each model required its own derivation of an exact CDF, its own boundary theorem, and its own EVT bridge, precisely because each new profile family's local anisotropy has a genuinely different functional form. What *has* been established, and is the more defensible general claim, is a repeatable method: identify the profile family's local anisotropy δ as a function of the sampled coordinates, identify the coarsest sampling model whose geometry matches that dependence, derive the exact CDF and boundary rate for that specific pairing, and only then ask whether a rescue statistic of κ_3's general type narrows the resulting bound. Section 32.1's quadratic-rate finding and the exponential-rate findings of Chapters 10 and 24 before it are three applications of this method to three profile families that happened to require three different closed forms, not three special cases that could have been anticipated from a single master theorem stated in advance. Whether a fourth profile family — Mollweide's genuinely two-dimensional shear, named as open immediately below, is the most concrete candidate — would require a fourth new closed form, or would turn out to reduce to one of the first three, is exactly the kind of question this method is built to answer by direct derivation rather than by extrapolation.

What remains genuinely open

Three items are named here, once, as concrete remaining directions rather than folded into the closing list above as though addressed by proxy.

1. **Mollweide's shear.** Section 29.3 showed the product reduction $Z = VW$ is specific

to Sinusoidal's $h^2 = 1 + \lambda^2 \sin^2 \phi$. Mollweide's local anisotropy away from its central meridian has a different closed form, not of this product type, and nothing in this Part establishes whether an analogous exact reduction exists for it; it may require a genuinely two-dimensional treatment with no such shortcut.

2. **The independence assumption.** Definition 29.1 samples latitude and longitude independently, a disclosed choice (Section 29.2) rather than a property of any real gazetteer. Whether Theorem 31.1's quadratic rate survives under a joint sampling density with genuine latitude–longitude correlation (as real place-name coverage has) is not tested here.

3. **The κ_3^{corner} bound-fit.** Chapters 12 and 26 each went on to test whether their model's rescue statistic improves a three-argument JSD bound fit against real GeoNames data. Section 29.3 disclosed at the outset that this Part does not attempt the analogous pipeline for κ_3^{corner}; doing so honestly would require a real two-margin coordinate sample, not a synthetic one, and is left as the concrete next step.

Chapter 33

Correction: the True Tissot Anisotropy

A correction, stated plainly

Chapters 30 and 31 used $\delta := h/k = \sqrt{1 + \lambda^2 \sin^2 \phi}$ as "the" local anisotropy of Sinusoidal away from its central meridian. This is the ratio of the two coordinate-axis scale factors, and every earlier chapter in this monograph (azimuthal cap model, latitude-band model) used exactly this ratio correctly, because in each of those cases the map's meridians and parallels remain orthogonal, so h/k *is* the true ratio of maximum to minimum local scale (Tissot's a'/b'). Sinusoidal's graticule is not orthogonal once $\lambda \neq 0$ and $\phi \neq 0$ — the very shear Chapter 29's title refers to — and h/k is not, in general, equal to a'/b' once the coordinate curves meet at an angle other than $90°$. This chapter derives a'/b' correctly, confirms the earlier ratio was not it, and rebuilds Chapters 30 and 31's results on the corrected quantity. The error is owned here rather than carried forward silently.

33.1 Tissot's indicatrix from the map's Jacobian

It is worth being precise about why h/k is the right quantity in every earlier chapter of this monograph and the wrong one here, rather than treating this as an isolated arithmetic slip local to Sinusoidal. h and k are, by definition, the scale factors *along the coordinate curves* — how much a small step in ϕ or in λ is stretched by the map, measured in each coordinate's own direction. When the coordinate curves cross at a right angle, as they do for every azimuthal profile and every cylindrical latitude-band profile used earlier in this cycle, those two coordinate directions already are the directions of maximum and minimum stretch, so h and k coincide with Tissot's true semi-axes a' and b' and h/k is exactly a'/b' with nothing further to derive. Sinusoidal's graticule loses that right angle away from the central meridian and equator — precisely the shear this Part's title names — and once the coordinate curves are oblique, the direction of maximum stretch no longer lies along either coordinate curve, so h and k measure stretch in the wrong two directions and h/k silently stops being a'/b'. The derivation below recovers a'/b' from first principles, via the map's Jacobian, precisely so that this dependence on orthogonality is never assumed again without being checked.

For a map $(\phi, \lambda) \mapsto (x, y)$, the sphere's orthonormal meridian/parallel frame is sent, at each point, to the plane vectors $\mathbf{v}_1 = (x_\phi, y_\phi)$ (image of the unit meridian direction) and $\mathbf{v}_2 = (x_\lambda, y_\lambda)/\cos\phi$ (image of the unit parallel direction, since a unit of arc length along a parallel corresponds to a λ-change of $1/\cos\phi$). These form the columns of a 2×2 matrix M, and Tissot's indicatrix at that point is the image of the unit circle under M: an ellipse whose semi-axes $a' \geq b' \geq 0$ are M's singular values — the true maximum and minimum local scale factors in *any* direction, not merely along the coordinate curves.

Proposition 33.1 (Semi-axes from trace and determinant)**.** *With* $h := \sqrt{x_\phi^2 + y_\phi^2}$, $k :=$

$\sqrt{x_\lambda^2 + y_\lambda^2}/\cos\phi$, *and area scale* $\sigma := |x_\phi y_\lambda - y_\phi x_\lambda|/\cos\phi$,

$$a'^2 + b'^2 = h^2 + k^2, \qquad a'b' = \sigma,$$
$$a' = \frac{\sqrt{h^2 + k^2 + 2\sigma} + \sqrt{h^2 + k^2 - 2\sigma}}{2}, \qquad b' = \frac{\sqrt{h^2 + k^2 + 2\sigma} - \sqrt{h^2 + k^2 - 2\sigma}}{2}. \tag{33.1}$$

Proof. $a'^2 + b'^2 = \mathrm{tr}(M^\top M)$ and $a'b' = |\det M|$, both immediate from the singular-value decomposition; $\mathrm{tr}(M^\top M) = h^2 + k^2$ by construction of M's columns, and $\det M = (x_\phi y_\lambda - y_\phi x_\lambda)/\cos\phi$ gives $|\det M| = \sigma$. Solving the two equations $a'^2 + b'^2 = h^2 + k^2$, $a'b' = \sigma$ for $a', b' \geq 0$ with $a' \geq b'$ gives Eq. (33.1). □

Where the earlier chapter went wrong

Chapter 30 implicitly assumed $a' = h$, $b' = k$ (equivalently $\sigma = hk$, i.e. orthogonal coordinate curves, $|\sin\theta'| = 1$ in the angle-between-curves sense). Section 33.2 below computes σ for Sinusoidal directly and shows $\sigma = 1 \neq hk = \sqrt{1 + \lambda^2 \sin^2\phi}$ whenever $\lambda \sin\phi \neq 0$: the coordinate curves are not orthogonal, and h/k therefore understates or misrepresents the true anisotropy a'/b'.

33.2 Sinusoidal's true anisotropy

For Sinusoidal ($x = \lambda\cos\phi$, $y = \phi$), direct differentiation gives $x_\phi = -\lambda\sin\phi$, $y_\phi = 1$, $x_\lambda = \cos\phi$, $y_\lambda = 0$, hence $h = \sqrt{1 + \lambda^2\sin^2\phi}$ (as before), $k = 1$ (as before), and area scale

$$\sigma = \frac{|x_\phi y_\lambda - y_\phi x_\lambda|}{\cos\phi} = \frac{|-\cos\phi|}{\cos\phi} = 1,$$

confirmed symbolically in `code/verify_sinusoidal_true_tissot_v2.py`: $\sigma \equiv 1$ identically, the exact expression of Sinusoidal's equal-area property (any true equal-area map has $\sigma \equiv 1$ in these units). With $h^2 + k^2 = 2 + u^2$, $u := |\lambda\sin\phi|$, Eq. (33.1) gives

$$a'(u) = \frac{u + \sqrt{u^2 + 4}}{2}, \qquad b'(u) = \frac{\sqrt{u^2 + 4} - u}{2}, \qquad a'(u)b'(u) = 1, \tag{33.2}$$

and the true local anisotropy is

$$\delta_{\text{true}}(u) := \frac{a'(u)}{b'(u)} = a'(u)^2 = \left(\frac{u + \sqrt{u^2 + 4}}{2}\right)^2, \qquad u = |\lambda\sin\phi|. \tag{33.3}$$

Equation (33.2) was checked against direct numerical SVD of M at four (ϕ, λ) points spanning the tested range, matching to six decimal places at every point (`code/verify_sinusoidal_true_tissot_v2.py`); Eq. (33.3) is not a re-derivation of Eq. (29.1) under a new name, but a materially different function of u: at small u, $\delta_{\text{true}}(u) \approx 1 + u + O(u^2)$ grows *linearly*, where the earlier $h/k = \sqrt{1 + u^2} \approx 1 + u^2/2$ grew quadratically and understated the true first-order distortion.

Key Result

Sinusoidal's true local anisotropy under shear is $\delta_{\text{true}}(u) = [(u+\sqrt{u^2+4})/2]^2$, $u = |\lambda \sin\phi|$, not $h/k = \sqrt{1+u^2}$ as Chapter 30 used. The two agree only at $u = 0$ (the undistorted central meridian/equator) and diverge from each other at first order in u elsewhere. Section 33.3 and Chapter 34 rebuild the exact CDF, boundary theorem, and EVT bridge on δ_{true}.

33.3 The corrected exact CDF

Section 29.3's reduction is unaffected by this correction: under Definition 29.1, $u = |\lambda \sin\phi| = aZ$ exactly as before ($a = \lambda_{\max} \sin x$, $Z = VW$), since the correction changes only the map from u to δ, not the sampling model or the reduction to Z. What changes is the inversion needed to state F_δ in terms of F_Z.

Proposition 33.2 (Corrected exact CDF)**.** *Let $u(t)$ invert Eq. (33.3), i.e. $t = [(u+\sqrt{u^2+4})/2]^2$. Then*

$$u(t) = \sqrt{t} - \frac{1}{\sqrt{t}}, \qquad t \in [1, D_{\text{true}}], \quad D_{\text{true}} := \delta_{\text{true}}(a). \tag{33.4}$$

Under Definition 29.1, $F_\delta(t) = F_Z\big(u(t)/a\big)$, with F_Z as in Eq. (30.1).

Proof. Write $s = \sqrt{t} = a'(u)$ (principal root, $t \geq 1 \Rightarrow s \geq 1$). From $s = (u+\sqrt{u^2+4})/2$: $2s - u = \sqrt{u^2+4}$, square both sides, $4s^2 - 4su + u^2 = u^2 + 4$, giving $s^2 - su = 1$, i.e. $u = s - 1/s = \sqrt{t} - 1/\sqrt{t}$. Composing with $u = aZ$ and inverting for Z gives $F_\delta(t) = \Pr[\delta \leq t] = \Pr[Z \leq u(t)/a] = F_Z(u(t)/a)$. □

Table 33.1 checks Proposition 33.2 against direct Monte Carlo simulation of δ_{true} under Definition 29.1 ($x = 0.30$, $\lambda_{\max} = 0.8$, so $a = 0.2367$, $D_{\text{true}} = 1.26601$), at six (p, q) combinations, eight test points t per cell, $n = 2 \times 10^6$.

Table 33.1: Maximum absolute deviation between empirical and exact corrected CDF (Eq. (33.4)), 8 test points $t \in [1.02, D_{\text{true}}]$, $n = 2 \times 10^6$. Expected MC noise scale ≈ 0.00071. Source: `code/corner_shear_true_tissot_cdf.py`, `data/corner_shear_true_tissot_cdf_results.json`.

p	q	$\alpha = 1/p$	Max |dev|
0.5	0.5	2.0	0.00049
1.0	1.0	1.0	0.00045
2.0	2.0	0.5	0.00013
0.5	2.0	2.0	0.00033
2.0	0.5	0.5	0.00056
1.0	2.0	1.0	0.00053

Every deviation is within the ≈ 0.00071 Monte Carlo noise band, confirming Proposition 33.2.

Key Result

The corrected exact CDF, Eq. (33.4) composed with Eq. (30.1), is confirmed numerically at six (p, q) combinations. Proposition 30.1's underlying machinery (the exact CDF of $Z = VW$) survives the correction untouched; only the map relating t to z needed to change, from $z = \sqrt{t^2-1}/a$ to $z = (\sqrt{t} - 1/\sqrt{t})/a$.

Two things about this repair are worth stating plainly, since a correction chapter invites a different kind of scrutiny than an ordinary results chapter. First, the scope of what needed fixing is narrower than it might first appear: Section 33.3's proof shows the entire product-CDF machinery of Proposition 30.1 — the harder half of Chapter 30's derivation, and the part independently verified against Monte Carlo at nine separate (p, q) cells there — required no change at all. The error was confined to the single link between the sampling model's Z and the geometric quantity δ, and once that link was corrected, everything built downstream of Z carried over intact. Second, the corrected numbers in Table 33.1 are not a weaker echo of Table 30.1's original verification: both tables show maximum deviations within a factor of two or three of the same ≈ 0.00071 noise floor, so the correction did not trade numerical precision for correctness. The chapters that follow build only on δ_{true}; Chapter 30's original h/k results are retained in this volume only as the record of what was corrected and why, not as a result still in use.

Chapter 34

The Corrected Boundary Theorem and EVT Bridge

34.1 The corner rate survives the correction; the constant does not

Section 31.1's mechanism — $F_Z'(1) = 0$, forcing $S_Z(1-\eta) \sim \frac{1}{2}\alpha\beta\eta^2$ — depended only on Definition 29.1's sampling structure, not on the specific map from Z to δ. Composed with any smooth, strictly increasing, invertible map G with $G'(a) \neq 0$ at the corner value $u = a$, the corner rate remains quadratic in ε; only the rate constant K, which depends on $G'(a)$, changes. This is confirmed below for $\delta_{\text{true}} = G(u)$, $G(u) = [(u + \sqrt{u^2+4})/2]^2$.

Theorem 34.1 (Corrected corner boundary rate)**.** *Under Definition 29.1, with true anisotropy δ_{true} as in Eq. (33.3), $\alpha = 1/p$, $\beta = 1/q$, $a = \lambda_{\max} \sin x$, $D_{\text{true}} = G(a)$,*

$$S(D_{\text{true}} - \varepsilon) \;\sim\; K\,\varepsilon^2 \quad \text{as } \varepsilon \to 0^+, \qquad K \;=\; \frac{\alpha\beta\,(a^2+4)}{8\,a^2\,D_{\text{true}}^2}. \tag{34.1}$$

Proof. Write $\delta_{\text{true}} = G(u)$, $u = aZ$. Near $Z = 1$ ($u = a$), $G(u) \approx D_{\text{true}} - a\,G'(a)\,\eta$ for $\eta = 1 - Z$, i.e. $\varepsilon \approx a\,G'(a)\,\eta$. From Eq. (33.2), $G'(u) = 2G(u)/\sqrt{u^2+4}$ (chain rule through $a'(u) = \sqrt{G(u)}$, using $a_u'(u) = a'(u)/\sqrt{u^2+4}$ and $G = a'^2$), so $G'(a) = 2D_{\text{true}}/\sqrt{a^2+4}$. Section 31.1's argument gives $S_Z(1-\eta) \sim \frac{1}{2}\alpha\beta\eta^2$; substituting $\eta \approx \varepsilon/(aG'(a)) = \varepsilon\sqrt{a^2+4}/(2aD_{\text{true}})$,

$$S(D_{\text{true}} - \varepsilon) = S_Z(1-\eta) \sim \frac{\alpha\beta}{2}\left(\frac{\sqrt{a^2+4}}{2aD_{\text{true}}}\right)^2 \varepsilon^2 = \frac{\alpha\beta(a^2+4)}{8a^2D_{\text{true}}^2}\,\varepsilon^2,$$

which is Eq. (34.1). □

34.2 Numerical verification

Table 34.1 reports the direct log–log slope fit (five geometrically-spaced ε values, $n = 2 \times 10^7$) at two $(x, \lambda_{\max}, p, q)$ combinations, together with the convergence of the empirical rate constant toward K_{theory}.

Table 34.1: Direct finite-ε verification of Theorem 34.1, $n = 2 \times 10^7$. Source: `code/corner_shear_true_tissot_cdf.py`, `data/corner_shear_true_tissot_boundary_results.json`.

$x, \lambda_{\max}, p, q$	D_{true}	K_{theory}	Fitted slope	$K_{\text{empirical}}$ (largest / smallest ε)	
$0.30, 0.8, 1, 1$	1.26601	5.659	1.985	5.963	5.685
$0.50, 1.2, 2, 1$	1.76413	0.263	1.974	0.274	0.267

Both fitted slopes $(1.985, 1.974)$ sit closer to the theoretical exponent 2 than Table 31.1's uncorrected fits did $(2.07–2.24)$, and $K_{\text{empirical}}$ brackets K_{theory} tightly as ε shrinks in both cases (relative error at the smallest well-populated ε under 5% in both cases; the very smallest ε in each table has a tail count in the low hundreds and correspondingly larger sampling noise, reported in full in `data/corner_shear_true_tissot_boundary_results.json`).

Key Result

Theorem 34.1's quadratic corner rate, with the corrected constant $K = \alpha\beta(a^2 + 4)/(8a^2 D_{\text{true}}^2)$, is confirmed directly. The qualitative finding of Chapter 31 — corner geometry gives a quadratic, not linear, boundary rate — was correct and survives the Tissot correction unchanged; only the specific constant changes, exactly as Section 34.1 predicted from the general chain-rule argument.

34.3 The corrected EVT bridge

Proposition 31.3's statement and proof carry over verbatim with D_{true} and the corrected K in place of the earlier D, K: $\sqrt{n}(D_{\text{true}} - M_n) \xrightarrow{d} T$, $\Pr[T > t] = \exp(-Kt^2)$. Table 34.2 reports the same two-route verification as Section 31.5: exact quadrature at three sample sizes per case, and an independent 4000-replicate Monte Carlo check.

Table 34.2: Exact-quadrature and Monte Carlo verification of the corrected EVT bridge. "Ratio" is quadrature (resp. MC) mean of $\sqrt{n}(D_{\text{true}} - M_n)$ over the theoretical Rayleigh mean $\sqrt{\pi/(4K)}$. Source: `code/corner_shear_true_tissot_evt.py`, `data/corner_shear_true_tissot_evt_results.json`, `data/corner_shear_true_tissot_evt_mc_results.json`.

$x, \lambda_{\max}, p, q$	K	n (quadrature)	Max |dev|	Ratio
$0.30, 0.8, 1, 1$	5.659	4,599	0.00482	0.9931
		18,399	0.00239	0.9965
		73,597	0.00119	0.9981
$0.50, 1.2, 2, 1$	0.263	12,005	0.00736	0.9899
		48,020	0.00367	0.9949
		192,083	0.00183	0.9974
Monte Carlo replicate check ($n = 2000$ samples, 4000 replicates):				
$0.30, 0.8, 1, 1$	5.659	—	—	0.9893
$0.50, 1.2, 2, 1$	0.263	—	—	0.9702

Both routes reproduce the same qualitative pattern found in Section 31.5: exact-quadrature deviation shrinking by close to a factor of 2 as n quadruples, quadrature mean ratio rising to within half a percent of 1 at the largest n tested, and an independent Monte Carlo check landing within 1–3% of the theoretical mean, the same small consistent finite-n/finite-replicate shortfall documented throughout this monograph.

Key Result

The corrected EVT bridge, $\sqrt{n}(D_{\text{true}} - M_n) \to \text{Weibull}(2)$ with the corrected rate K, is confirmed by exact quadrature and independent Monte Carlo. Chapter 29's open item — whether the boundary/EVT machinery extends to a genuinely two-variable, shear-bearing profile — is now closed on the *correctly defined* local anisotropy, not the coordinate-ratio approximation Chapters 30 and 31 used. Chapter 35 takes up the two remaining named

items from Chapter 32: whether this reduction extends to Mollweide, and what the independence assumption in Definition 29.1 does and does not cost.

34.4 The convergence rates in full

Two separate convergence claims are made in this chapter, and it is worth stating both to the same numerical precision used elsewhere in this monograph rather than only in the qualitative terms of the two keyresult boxes above. First, the rate-constant convergence in Table 34.1: at $(x, \lambda_{\max}, p, q) = (0.30, 0.8, 1, 1)$, $K_{\text{empirical}}$ moves from 5.963 at the largest tested ε to 5.685 at the smallest, against $K_{\text{theory}} = 5.659$ — a relative error falling from 5.4% to 0.5%; at $(0.50, 1.2, 2, 1)$, the corresponding move is from 0.274 to 0.267 against $K_{\text{theory}} = 0.263$, a relative error falling from 4.2% to 1.5%. Both cases show the error shrinking monotonically as $\varepsilon \to 0$, which is the specific signature Section 34.2 calls confirmation, as distinct from a fitted constant that merely happens to lie near the theoretical value at one scale but does not tighten as the asymptotic regime is approached.

Second, the exact-quadrature convergence in Table 34.2: the maximum deviation at $(0.30, 0.8, 1, 1)$ runs $0.00482 \to 0.00239 \to 0.00119$ as n quadruples $(4599 \to 18399 \to 73597)$, a ratio of 2.02 then 2.01 between successive steps — matching the $\sqrt{n}$ scaling the theorem predicts (quadrupling n should halve a $1/\sqrt{n}$-type deviation) to within one part in a hundred at both steps, not merely in the right direction. The second case, $(0.50, 1.2, 2, 1)$, shows the same pattern: $0.00736 \to 0.00367 \to 0.00183$, ratios of 2.01 and 2.01. This internal consistency — the deviation shrinking at almost exactly the rate the asymptotic theory demands, rather than at some other rate that happens to also decrease — is a stronger check than the raw deviation figures alone convey, and is the reason Section 34.3 treats the quadrature route as confirmatory rather than merely suggestive, with the independent 4000-replicate Monte Carlo check (landing at 0.9893 and 0.9702 of the theoretical mean in the two cases) serving as an independent cross-check using a different computational method entirely, rather than a repetition of the same quadrature procedure at different settings.

Chapter 35

Mollweide, and What Independence Costs

35.1 Why Mollweide does not inherit the reduction

Mollweide's defining equations, $x = \frac{2\sqrt{2}}{\pi}\lambda\cos\theta$, $y = \sqrt{2}\sin\theta$, with $\theta(\phi)$ solved from $2\theta + \sin 2\theta = \pi\sin\phi$, share Sinusoidal's structural feature that y does not depend on λ and x depends on λ only linearly; Proposition 33.1's machinery applies without modification. Writing $h^2 = A(\phi) + B(\phi)\lambda^2$ (only the x_ϕ term carries λ, exactly as for Sinusoidal) and $\sigma = C(\phi)$ (the area scale, which Proposition 33.1's $\det M$ shows depends on ϕ alone here too, since $y_\lambda = 0$ kills any λ-dependence in the Jacobian just as it did for Sinusoidal), Eq. (33.1) gives

$$(a'+b')^2 = Q_+(\phi)+B(\phi)\lambda^2, \qquad (a'-b')^2 = Q_-(\phi)+B(\phi)\lambda^2, \qquad Q_\pm(\phi) := A(\phi)+k(\phi)^2 \pm 2C(\phi). \tag{35.1}$$

Sinusoidal's clean reduction (Eq. (33.2)) is the special case $Q_-(\phi) \equiv 0$: there, $(a' - b')^2 = B\lambda^2$ exactly, so $a' - b'$ is proportional to $|\lambda|\sqrt{B(\phi)} = u$ alone, collapsing two variables into one combined variable cleanly. Table 35.1 computes $Q_-(\phi)$ for Mollweide directly (numerically solving $\theta(\phi)$ by root-finding, differentiating numerically for $\theta'(\phi)$, both in `code/mollweide_reduction_check.py`) at seven latitudes.

Table 35.1: Mollweide's $Q_-(\phi)$ (Eq. (35.1)), at seven latitudes, computed via numerical root-finding for $\theta(\phi)$. Sinusoidal's corresponding quantity is 0 identically at every ϕ (Section 33.2). Source: `code/mollweide_reduction_check.py`.

ϕ (deg.)	5	15	30	45	60	75	85
$Q_-(\phi)$	0.0430	0.0338	0.0101	0.0027	0.0943	0.6283	2.7845

What this means, plainly

$Q_-(\phi) \neq 0$, and grows sharply toward the poles, for Mollweide at every latitude tested. Sinusoidal's reduction to a single combined variable $u = |\lambda\sin\phi|$ is therefore not a generic feature of equal-area pseudocylindrical projections; it is a special consequence of Sinusoidal's particular closed form (specifically $\sigma \equiv 1$ *and* h^2 depending on λ, ϕ only through the product $\lambda\sin\phi$, which together force $Q_- \equiv 0$). For Mollweide, a' and b' genuinely depend on ϕ and λ separately, not through a single combined variable, and Chapters 30 and 34's reduction to a one-dimensional exact CDF in $Z = VW$ does not apply. A Mollweide boundary/EVT treatment would need either a genuinely two-dimensional distributional argument, or a numerical (Monte-Carlo-only, no closed-form CDF, since $\theta(\phi)$ itself has no closed form) treatment built from scratch. Neither is attempted here: this is named as the concrete next chapter this Part does not contain, not folded into the results above as though it were covered by them.

35.2 The shape of $Q_{-}(\phi)$, and why it is not monotone

Table 35.1's seven values are worth reading as a shape, not only as a list of nonzero numbers, since the shape itself carries information about where a Mollweide boundary/EVT treatment would be hardest to build. $Q_{-}(\phi)$ does not grow monotonically from the equator to the pole: it falls from 0.0430 at $5°$ to a minimum of 0.0027 at $45°$ — roughly a sixteenfold decrease — before rising sharply to 0.0943 at $60°$, 0.6283 at $75°$, and 2.7845 at $85°$. The ratio between the largest and smallest tabulated values, $2.7845/0.0027 \approx 1031$, spans three orders of magnitude across the tested latitude range, and the location of the minimum near $45°$ rather than at the equator is itself informative: it shows the near-equatorial values are not simply the smallest case of a single monotone trend that could be linearly extrapolated from low latitudes, and any future numerical treatment of Mollweide's two-dimensional shear would need to resolve this non-monotone shape directly rather than assume $Q_{-}(\phi)$ is small and roughly constant away from the poles.

This pattern is also a useful diagnostic against the possibility that the nonzero $Q_{-}(\phi)$ values in Table 35.1 are a numerical artefact of the root-finding and finite-difference procedure in `code/mollweide_reduction_check.py` rather than a genuine geometric feature. An artefact of numerical differentiation would typically appear as noise of roughly constant magnitude across ϕ, or would grow smoothly with the step size used near the pole where $\theta(\phi)$'s implicit equation becomes harder to solve; a value that falls by a factor of sixteen and then rises by three orders of magnitude, tracking a specific and reproducible functional shape rather than a noise floor, is the signature of a real effect being measured, not of accumulating solver error. The steep late rise is also consistent with $Q_{-}(\phi)$'s definition in Eq. (35.1): near the poles, Mollweide's meridians converge toward the projection's pole-points while $\theta(\phi)$'s implicit solution becomes increasingly sensitive to ϕ, which is exactly the regime in which $A(\phi)$, $C(\phi)$, and $k(\phi)$ would be expected to separate most sharply from the near-cancellation that keeps Q_{-} small in the middle latitudes.

35.3 What the independence assumption in Definition 29.1 does and does not cost

Section 29.2's disclaimer flagged sampling ϕ and λ independently as a disclosed simplification. One part of that concern can be settled directly, without new data.

Proposition 35.1 (Independence is exact for area-true sampling of a rectangular window)**.** *For a rectangular latitude–longitude window $\phi \in [0, x]$, $\lambda \in [-\lambda_{\max}, \lambda_{\max}]$, sampling points uniformly with respect to true surface area (i.e. with density proportional to $\cos\phi$, the standard area-true convention this monograph has used throughout, e.g. Definition 23.1) draws ϕ and λ* independently, *because the area element $dA = \cos\phi\, d\phi\, d\lambda$ factors exactly into a function of ϕ alone times a function of λ alone (the constant 1).*

Proof. The joint density restricted to the window is $f(\phi, \lambda) \propto \cos\phi \cdot \mathbf{1}_{[0,x]}(\phi) \cdot \mathbf{1}_{[-\lambda_{\max},\lambda_{\max}]}(\lambda)$, a product of a function of ϕ alone and a function of λ alone; a jointly distributed pair with product density is independent by definition. □

So what is actually still open

Proposition 35.1 means Definition 29.1's independence at $p = q = 1$ is not an approximation at all for a rectangular window sampled area-true: it is the exact consequence

of dA factoring. The concentration exponents $p, q \neq 1$ used elsewhere in this Part are a further disclosed modeling choice (concentrating samples toward the window's edges or centre, independently in each margin), unaffected either way by this proposition. What Proposition 35.1 does *not* settle is whether a real gazetteer's place-name coverage over a real geographic region resembles uniform (or p, q-concentrated) sampling over a rectangular window at all: real coverage correlates with population density, coastlines, and administrative boundaries in ways that are not simply a product of a latitude margin and a longitude margin, and no amount of reasoning about the idealized rectangular model settles that empirical question. That remains open, and is exactly the question Section 36.2 returns to.

Chapter 36

Part Four: Final Closing Accounting

This chapter supersedes Chapter 32's accounting

Chapter 32 closed this Part on the uncorrected $\delta = h/k$. Chapter 33 found and fixed that error. This chapter restates the accounting on the corrected results, and updates the open-items list accordingly; Chapter 32 is left in place as part of the honest record of how this Part actually proceeded, not deleted.

36.1 What is closed

- **The two-variable profile itself** (Chapter 29): Sinusoidal's scale factors away from the central meridian, symbolically re-derived and confirmed.
- **The sampling model** (Definition 29.1): an explicit, disclosed two-margin model, and (Proposition 35.1) a proof that its independence is in fact exact, not approximate, for area-true sampling of a rectangular window — one of Chapter 32's three named open items is fully resolved, not merely narrowed.
- **The exact CDF, corrected** (Chapter 33): $F_\delta(t) = F_Z((\sqrt{t} - 1/\sqrt{t})/a)$ on the true Tissot anisotropy $\delta_{\text{true}} = [(u + \sqrt{u^2+4})/2]^2$, replacing the earlier coordinate-ratio approximation, confirmed to Monte Carlo noise at six (p, q) combinations.
- **The boundary theorem and EVT bridge, corrected** (Chapter 34): the corner-type quadratic rate and $\sqrt{n}$-rescaled Weibull(2) limit both survive the correction in kind (confirming Section 34.1's general chain-rule argument that corner geometry, not the specific profile formula, drives the quadratic exponent), with corrected constants confirmed by direct slope fit, exact quadrature at three sample sizes, and independent Monte Carlo replication.
- **Why Mollweide does not inherit the reduction** (Chapter 35): not merely asserted but shown, via the same Tissot machinery applied to Mollweide's actual (numerically-solved) $\theta(\phi)$, that its analogue of $Q_-(\phi)$ is nonzero and latitude-dependent, so the one-variable collapse is a genuine special feature of Sinusoidal's closed form and not a general property of shear-bearing pseudocylindricals.

36.2 What remains open

1. **A Mollweide boundary/EVT treatment.** Section 35.1 showed the existing reduction does not apply; a treatment would need a genuinely two-dimensional argument or a purely numerical one (no closed-form CDF is available, since $\theta(\phi)$ itself is defined only implicitly), and is not attempted in this Part.

2. **Real, non-rectangular, correlated coverage.** Proposition 35.1 settled the idealized rectangular-window case exactly; whether real gazetteer coverage over an actual geographic region resembles it is an empirical question this Part's synthetic sampling model cannot answer on its own.

3. **The κ_3^{corner} bound-fit against real data.** At the time this chapter was first written, this was the item this session was least able to close: the tool environment's network access is restricted to a fixed allow-list of package and source-code repositories, and does not include `geonames.org` or any other gazetteer host. A real GeoNames extract was subsequently supplied directly, and Chapter 37 builds and runs the pipeline this item called for. It is left listed here, unedited, as part of the honest record of what was and was not reachable at each point in this session; Chapter 37 is the resolution, not a silent retroactive fix to this paragraph.

On the correction itself

Chapter 33 is included in this build not as a footnote but as a full chapter, because the error it fixes was substantive: h/k and the true a'/b' diverge at first order in the shear variable u, not merely in a higher-order correction. The same no-invented-data discipline that produced the original (incomplete) result is what caught it — checking the claimed anisotropy against an independent SVD computation of the actual Tissot ellipse, rather than trusting the coordinate-ratio formula by analogy with the orthogonal-grid profiles studied earlier in this monograph.

36.3 How large the correction was, and how tightly it was re-verified

It is worth stating in one place how the corrected model was checked, since "corrected" is a claim that itself needs the same standard of evidence as any other result in this monograph, not an exemption from it. The corrected exact CDF, $F_\delta(t) = F_Z\big((\sqrt{t} - 1/\sqrt{t})/a\big)$ on the true Tissot anisotropy δ_{true}, was tested against direct Monte Carlo simulation at six (p, q) combinations spanning both symmetric margins ($p = q \in \{0.5, 1, 2\}$) and asymmetric ones ($p \neq q$, tested in both directions). The maximum absolute deviation between the exact CDF and its Monte Carlo estimate across all nine tested quantiles was, at worst, 5.6×10^{-4} (at $p = 2, q = 0.5$), and at best 1.3×10^{-4} (at $p = q = 2$); every one of the six cases fell within Monte Carlo noise for the sample sizes used, with no systematic trend across the (p, q) grid that would indicate a residual, uncorrected shear term. This is the numerical content behind the phrase "confirmed to Monte Carlo noise" in the closed-items list above.

It is also worth being precise about what the correction did and did not change, since the disclaimer box above states the qualitative survival of the boundary theorem and EVT bridge but not the quantitative cost of getting there. Section 34.1's chain-rule argument shows that the quadratic boundary rate and the Weibull(2) limit law are consequences of the corner geometry — two margins entering multiplicatively near a joint extreme — rather than of the specific algebraic form of h/k versus δ_{true}; correcting the anisotropy formula therefore changes the constants entering Theorem 31.1 and Proposition 31.3 but not their functional type. This is precisely the sense in which Chapter 32 was wrong in the numbers it reported but right in the structural claim it made, and is the reason this chapter supersedes rather than retracts it.

Finally, on the open item concerning real, correlated coverage: Proposition 35.1's proof that the two-margin sampling model's independence is exact — not merely assumed — for area-true sampling of a rectangular window is a stronger result than a synthetic model normally earns,

precisely because it was checked rather than taken as the model's premise. It narrows, without eliminating, the second open item listed above: the rectangular-window case no longer needs the independence assumption disclaimed elsewhere in this monograph, and what remains open is only whether real gazetteer coverage over an actual, non-rectangular geographic region approximates that rectangular idealization closely enough for the quadratic rate to survive in practice — precisely the question Chapter 37 takes up next, using the GeoNames extract subsequently supplied for that purpose.

Chapter 37

A Real-Data Check: Georgia and France

The first real point cloud in this volume

Every earlier chapter of this synthesis volume — Chapters 6, 12 and 26 included — used synthetic sampling models, and each one named a real-data check as a natural next step it was not taking (Sections 5.2 and 26.2's closing remarks, Chapter 36's third open item). Real GeoNames extracts for twenty-one countries, supplied directly for this purpose, make that check possible for the first time in this volume. This chapter builds the corner-shear model's own JSD bound-fit (paralleling Chapters 12 and 26) and then, going one step further than any earlier chapter in this volume, tests it against two real countries' actual populated-place coordinates.

37.1 The synthetic corner-shear bound-fit

Following Section 26.1's pipeline exactly, with the corner-shear sampling model (Definition 29.1) in place of the band model and Sinusoidal coordinates $(X, Y) = (\lambda \cos\phi, \phi)$ in place of Equirectangular's: sample points, project, compute the JSD between pairwise geodesic and pairwise planar distance histograms, fit $\log \text{JSD} = a \log(\kappa - 1) + b \log(D - 1) + d \log \kappa_3^{\text{corner}} + c$. Two scope choices are made explicit here, both for tractability: $\lambda_{\max}$ is held fixed at 0.8 rather than swept as a fourth grid dimension, and $q = p$ (the two concentration exponents are tied together) rather than swept independently. The grid is $x \in \{0.15, 0.30, 0.50\}$, $p = q \in \{0.2, 0.35, 0.5, 0.7, 1.0, 1.5, 2.0, 3.0, 5.0\}$, $N = 1500$ points, 8 replicates per cell, $p \leq 3.0$ used for the fit (24 cells) and $p = 5.0$ held out as an out-of-range test (3 cells) — the same split Chapters 12 and 26 used.

Table 37.1: Two-argument versus three-argument (κ_3^{corner}) fit, corner-shear model, Sinusoidal, 24-point grid, $p \leq 3.0$. Source: `code/corner_kappa3_bound.py`, `data/corner_kappa3_bound_results.json`.

Fit	R^2	In-range span	Condition number
Two-argument (κ, D)	0.7537	$47.8\times$	—
Three-argument (κ, D, κ_3^{corner})	0.9599	$6.0\times$	21.5

A wide, noisy dynamic range

Unlike the band model's narrow JSD range (Table 26.2), the corner-shear grid's JSD spans roughly six orders of magnitude (4.9×10^{-7} to 3.0×10^{-3} across the 27 cells), closer to the cap model's wide range than the band model's narrow one. The smallest cells (gentlest sampling, $x = 0.15$, high $p = q$) have relative Monte Carlo noise of several tens of percent on the JSD estimate itself (`data/corner_kappa3_bound_results.json` records each cell's std alongside its mean); the fit above is reported as executed, on this genuinely noisy grid,

not smoothed or re-run to look tighter.

Adding κ_3^{corner} raises R^2 from 0.75 to 0.96 and narrows the in-range span nearly eightfold ($47.8\times \rightarrow 6.0\times$), and at the $p = 5$ out-of-range test the three-argument fit improves every cell (two-argument ratios 73–126×; three-argument ratios 1.7–3.3×) — in range and moderately out of range, this reproduces Chapter 12's cap-model pattern (a real, large rescue) rather than Chapter 26's modest one, consistent with this grid's wide dynamic range giving κ_3^{corner} more genuine work to do.

37.2 Two real countries

Georgia (GE, 5,419 populated places) and France (FR, 81,598) were each subsampled to $N = 1{,}500$ (matching the synthetic grid's point count), projected with Sinusoidal centred on each country's own mean longitude, and run through the identical JSD pipeline. D_{real} uses each country's actual bounding extent (maximum absolute latitude and maximum longitude offset from its own central meridian, over the *full* real dataset, not the subsample) in place of the synthetic model's $(x, \lambda_{\max})$.

Table 37.2: Real-data results, Georgia and France, $N = 1500$ populated places each. Source: `code/real_country_corner_check.py` (Georgia), `code/real_country_corner_check_FR.py` (France); `data/real_country_corner_check_results.json`, `data/real_country_corner_check_FR_results.json`.

Quantity	Georgia	France
Central meridian	43.33°E	2.36°E
Real extent: $\lvert\phi\rvert_{\max}$, $\lvert\lambda\rvert_{\max}$	43.58°, 3.33°	51.08°, 7.49°
D_{real}	1.0409	1.1070
κ_{real}	1.0137	1.0290
κ_3^{corner} (real)	0.0487	0.1843
JSD (real, Sinusoidal)	1.0×10^{-6}	1.0×10^{-5}
Two-argument prediction	3.0×10^{-7}	1.7×10^{-6}
Two-argument ratio	3.33×	5.77×
Three-argument prediction	3.2×10^{-9}	8.3×10^{-9}
Three-argument ratio	317×	1202×

A finding worth stating plainly: the rescue reverses

Both real countries' D_{real} falls *below* the synthetic grid's tested range ($[1.1269, 1.4641]$): real national territories, even ones spanning 8–10 degrees of latitude at temperate latitudes, have D_{real} gentler than the mildest synthetic cell tested. This is itself informative (Section 37.3 explains why), but the sharper finding is what happens to the two fits when extrapolated into this untested, gentler regime: the two-argument fit, while imperfect (3.3–5.8× over the real JSD), stays within an order of magnitude; the three-argument fit — which was the clear winner in range and at the moderately-out-of-range $p = 5$ test above — degrades catastrophically (317× and 1202×). This is the opposite of every earlier in-range/out-of-range comparison in this monograph, where adding κ_3 (or κ_3^{corner}) never made an out-of-range prediction worse than the two-argument bound it was built to improve. It is reported here exactly as found: a genuine limitation of the three-argument fit's extrapolation behaviour, not smoothed over because it complicates an otherwise clean rescue story. Section 37.4 takes this reversal as a diagnostic question rather than a closed

one, and finds it is substantially, though not fully, a range-of-calibration artifact.

37.3 Why real countries sit in a gentler regime

The corner-shear model's rate constant depends on $a = \lambda_{\max} \sin x$ (or, for real data, on the actual product of longitude half-extent and sin of latitude extent), and both real countries tested have a longitude half-extent under $7.5°$ (0.13 rad), against the synthetic grid's fixed $\lambda_{\max} = 0.8$ rad ($45.8°$) — a more than sixfold difference that dominates a regardless of latitude. This reflects a genuine geographic regularity the synthetic grid was not built to capture: national territories, even large ones spanning many degrees of latitude, are typically far narrower in longitude than they are tall in latitude (France's $9.3°$ latitude span against a $14.9°$ longitude span is one of the wider real ratios among typical countries; Georgia's is far more elongated, $2.4°$ latitude against $6.7°$ longitude, in the other direction). Definition 29.1's $\lambda_{\max} = 0.8$ was chosen in Section 37.1 to give the corner effect enough room to be worth testing at all, not to represent a typical country's actual shape — a scope choice now shown, by the real-data check this chapter set out to run, to sit well outside typical real geography.

37.4 Extending the grid: does an explicit $\lambda_{\max}$ term fix it?

Section 37.1's fit held $\lambda_{\max}$ fixed at 0.8 and let κ, D, and κ_3^{corner} absorb its effect implicitly. But κ and D both depend on x and $\lambda_{\max}$ chiefly through the product $a = \lambda_{\max} \sin x$ (Section 37.3): different $(x, \lambda_{\max})$ pairs can share the same D while producing point clouds of very different shape — one tall and narrow, the other short and wide. Since JSD depends on the full pairwise-distance-histogram, not just on the boundary value D, this is a genuine geometric degeneracy that κ, D, κ_3^{corner} alone need not resolve. Two questions follow directly from Section 37.2's reversal: does re-fitting on a grid that actually spans real countries' much smaller $\lambda_{\max}$ fix the extrapolation, and does adding $\log \lambda_{\max}$ as an explicit fourth regressor fix it further?

A disclosed reduction in grid resolution

The extended grid multiplies the (x, p, q) grid by five $\lambda_{\max}$ values (135 cells total against Section 37.1's 27), which was not tractable at the original $N = 1500$ points and 8 replicates per cell within this session. N was reduced to 1200 and replicates to 5; this is a disclosed scope choice made for computational tractability, not a change intended to flatter the result, and it means the extended-grid numbers below are not directly comparable cell-by-cell to Table 37.1's in precision, only in the pattern they show.

The extended grid uses $x \in \{0.15, 0.30, 0.50\}$ (unchanged) and $\lambda_{\max} \in \{0.05, 0.10, 0.20, 0.40, 0.80\}$, bracketing both real countries' actual $\lambda_{\max}$ (Georgia 0.0582 rad, France 0.1307 rad) and Section 37.1's original fixed value. $p = q$ uses the same nine-value grid, $p \le 3.0$ fit (120 cells) and $p = 5.0$ held out (15 cells), the same split convention as before.

Against Table 37.1's narrow-grid three-argument fit ($R^2 = 0.9599$, span $6.0\times$), both extended-grid fits are markedly worse in range: R^2 falls to roughly 0.75 and the span widens nearly fortyfold, to 244–$271\times$. Adding $\log \lambda_{\max}$ explicitly barely moves R^2 ($0.7538 \to 0.7582$) and, if anything, widens the span slightly while roughly doubling the condition number ($29.9 \to 39.6$, reflecting near-collinearity between $\log \lambda_{\max}$ and the other three regressors). At the held-out $p = 5.0$ cells, the pattern repeats: ratios range 0.44–$56.6\times$ (span $128\times$) for the three-argument extended fit and 0.56–$60.2\times$ (span $107\times$) for the four-argument one — both far looser than the

Table 37.3: Extended grid ($\lambda_{\max} \in \{0.05, 0.10, 0.20, 0.40, 0.80\}$, $x \in \{0.15, 0.30, 0.50\}$, $p = q \leq 3.0$, 120 cells, $N = 1200$, 5 reps): the same three-argument fit re-run on the wider grid, versus a fourth regressor $\log \lambda_{\max}$ added explicitly. Source: `code/corner_kappa3_lammax_extended.py`, `data/corner_kappa3_lammax_extended_results.json`.

Fit	R^2	In-range span	Condition number
Three-argument (κ, D, κ_3^{corner}), extended grid	0.7538	244.0×	29.9
Four-argument ($+\log \lambda_{\max}$), extended grid	0.7582	271.0×	39.6

narrow grid's 1.7–3.3×. A single additional log-linear term does not resolve the degeneracy diagnosed above.

Table 37.4: Real-country extrapolation ratios (actual JSD / predicted JSD): narrow grid ($\lambda_{\max} = 0.8$ only, Table 37.2) versus the extended grid of this section. The two-argument fit is unaffected by this section's changes and repeated for reference.

Fit	Georgia	France
Two-argument (κ, D), narrow grid	3.33×	5.77×
Three-argument (κ, D, κ_3^{corner}), narrow grid	317×	1202×
Three-argument, extended grid	1.86×	8.32×
Four-argument ($+\log \lambda_{\max}$), extended grid	1.38×	6.31×

Here the picture reverses again, and favourably: merely re-fitting the unchanged three-argument model on a grid that spans real countries' actual $\lambda_{\max}$ range turns the catastrophic narrow-grid failure (317×, 1202×) into ratios (1.86×, 8.32×) that are comparable to, or for Georgia better than, the plain two-argument fit's. Adding $\log \lambda_{\max}$ explicitly improves both further (1.86× → 1.38×; 8.32× → 6.31×). So the original reversal was substantially, though evidently not entirely, an artifact of calibrating at a single $\lambda_{\max} = 0.8$ and extrapolating to real countries' much smaller values — not a fundamental defect in κ_3^{corner} itself.

A rescue that is real but only partial

Both results in this section are genuine and hold together without contradiction: extending the calibration range fixes most of the extrapolation failure (two-to-three orders of magnitude down to low single digits or, for France, a moderate one order), while simultaneously making the in-range fit itself substantially noisier (R^2 0.96 → 0.75, span 6× → 244–271×). Both follow from the same degeneracy: once $\lambda_{\max}$ varies independently of x, κ and D no longer pin down the point cloud's shape tightly enough for a clean power-law fit, whether or not $\lambda_{\max}$ is added as an explicit term. Widening the calibration range so a fit reaches real data is not the same as finding a functional form that fits well *and* reaches real data — this section obtained the first without the second.

Key Result

The real-data check named as open in Chapter 36 is now done, for two countries of different shape and size, and the further diagnostic step this chapter's original ending named as undetermined — whether the reversal is a range artifact or a deeper instability — has now been run rather than left open. The corner-shear model's boundary theorem and EVT bridge (Chapter 34) rest on sound, independently-verified mathematics; the narrow-grid κ_3^{corner} term is a real, substantial in-range and moderately-out-of-range improvement over

κ, D alone; and its catastrophic real-country extrapolation failure $(317\times,\ 1202\times)$ is now shown to be substantially a range-of-calibration artifact, not a fundamental defect in the term itself: re-fitting on a $\lambda_{\max}$-extended grid brings both real countries to within a single order of magnitude (1.4–$8.3\times$ with $\lambda_{\max}$ added explicitly), comparable to the plain two-argument bound. The cost is real too, and is reported alongside the improvement rather than in place of it: the extended grid's own in-range fit is markedly looser ($R^2\ 0.96 \to 0.75$, span $6\times \to 244$–$271\times$), and a single additional log-linear regressor does not close that gap. What remains open, now narrower and better characterised than before, is whether a non-log-linear functional form — built from x and $\lambda_{\max}$ directly, or from an aspect-ratio quantity such as $x/\lambda_{\max}$, rather than through κ, D, κ_3^{corner}'s derived combinations — can recover the narrow grid's tight in-range fit while keeping this section's good extrapolation. That is a distinct modelling exercise, not another regression on the same three features, and is named here as the concrete next step.

References

Bibliography

[Salvador, Vol. I] Salvador, S. *Original Mathematical Perspectives on Cartography*.

[Salvador, Vol. II] Salvador, S. *Empirical Informational Distortion of Map Projections: A Cross-Regional Study on Real Settlement Distributions*. Companion volume to Vol. I.

[Salvador, Vol. III] Salvador, S. *Metric Embeddings, Impossibility and Information Geometry in Cartography: Distortion, Limits, and the Statistical Geometry of Map Projections*. Independent companion to Vol. I.

[Snyder, 1987] Snyder, J. P. *Map Projections: A Working Manual*. U.S. Geological Survey Professional Paper 1395, 1987.